9

Sixth Semiannual Technical Report. no. 6,

February 1976

11

12 94p.

For the Project

6

Local, Regional and Large Scale Integrated Networks.

VOLUME 1.

LARGE NETWORKS FOR DEFENSE COMMUNICATIONS: STRATEGIES, PERFORMANCE MEASURES AND COST TRENDS.

10

Principal Investigator: Howard Frank,

Co-principal Investigators: Israel Gitman
Richard Van Slyke

Contractor

NETWORK ANALYSIS CORPORATION
Beechwood, Old Tappan Road
Glen Cove, NY 11542
(516) 671-9580

ARPA Order No. 2286

Contract No. DAHC 15-73-C-0135

15

Effective Date: 13 October 1972
Expiration Date: 30 June 1977

Sponsored by

Advanced Research Projects Agency
Department of Defense

1473

389 161 LB

SIXTH SEMIANNUAL TECHNICAL REPORT

TABLE OF CONTENTS

VOLUME 4

LOCAL AND REGIONAL COMMUNICATION NETWORKS - TECHNOLOGIES AND ARCHITECTURES

VOLUME 5

COMPUTATIONAL TECHNIQUES

SUMMARY OF VOLUME 1

TECHNICAL PROBLEM

Network Analysis Corporation's contract with the Advanced Research Projects Agency in the area of Local, Regional, and Large Scale Integrated Networks, has the following objectives:

- Determine the feasibility and range of applicability of existing and developing communication technologies, for local, regional, and large scale data communication networks.

- Contribute to the development of new communication technologies by developing algorithms for network operation (e.g., routing, flow control), evaluating hardware and software alternatives, and assessment of prototype communication systems.

- Develop algorithms and efficient computational techniques for minimum cost design, for traffic analysis, and for reliability analysis of local, regional, and large scale data communication networks.

- Apply computational techniques to DOD data communication requirements, to determine the multi-dimensional cost/performance tradeoffs for each of the technologies under consideration.

GENERAL METHODOLOGY

The approach to the solution of these problems has been the simultaneous,

- Study of fundamental network analysis and design issues.
- Development of efficient algorithms for large scale network analysis and design.
- Development of an interactive distributed display and computational system to deal with large-scale problems.
- Application of the new analysis and design tools to study cost and performance tradeoffs for large systems.

TECHNICAL RESULTS

The detailed technical results are presented in Volumes 2 through 5 of this report. This volume integrates and summarizes the results discussed in the remaining volumes of this report. Emphasis here is placed on problem formulation, issues, and conclusions, rather than on the techniques used to reach these conclusions. The major conclusions summarized in this volume are listed below:

Network Structure

- In a large well-designed network, the most cost-effective network architecture is a multi-level hierarchical structure. It contains a

high level backbone (long haul) network and low level local access networks. Both the high and low level networks may be hierarchical. For example, the backbone network may contain satellite and terrestrial subnetworks.

- In a large optimized hierarchical network, delay and reliability in the backbone network are typically one order of magnitude better than in the local distribution network (except for satellite delays) and therefore variations in backbone delay and reliability requirements have little effect on global performance.

Local Access Networks

- Local access costs are on the order of 50% (or more) of total network costs. Moreover, total system cost is most sensitive to changes in local access component cost variations. Therefore, the local distribution network should receive high priority for cost minimization.

- Switch hardware and backbone communication line costs share the remaining cost on about an equal basis.

- The most cost-effective conventional local distribution alternative is obtained by using dedicated point-to-point lines which achieve cost savings via clustering of colocated terminals and Hosts by means of multiplexing or concentration.

- Multidrop polling techniques on tree or loop structures do not offer sufficient savings to warrent the changes required to overcome implementation difficulties and security shortcomings. This is significant because it implies that no change of strategy is needed for cost consideraations.

- Satellite use for local access via ground stations located at terminal or Host facilities does not appear to be cost-effective at current ground station costs and kilobit data rate requirements per facility.

- Packet radio provides a cost-effective alternative to conventional local access techniques. Packet Radio Systems can meet the high reliability requirements for critical DOD systems. If PRU's were manufactured in large quantities for $20,000 - $30,000 or less, PRNET designs would be more cost-effective than conventional designs. These costs are in the range of current technology and trends indicate that these units can be built for substantially less than this cost.

Long Haul Backbone Networks

- Local access and backbone network cost and performance is interrelated by the number, location, and capacity of the backbone switching nodes. Total system cost versus the number of backbone nodes does not vary significantly over a wide range of backbone node numbers, and hence, the

decision to distribute switches to many locations or consolidate switching as in AUTODIN II - Phase I is relatively independent of communication network economics.

- The optimum number of backbone nodes for AUTODIN II - Phase I ranges from 9 to 12 sites provided that switching facilities have sufficient capacity to handle traffic requirements. However, cost differentials are less than 10% when the number of backbone switches are increased substantially. Hence, wide distribution of switches, a desirable survivability feature, may be achieved at small incremental hardware and line cost.

- Considerable backbone line cost savings (on the order of 50%) can be achieved by installing satellite ground stations at all <u>backbone</u> nodes and eliminating all terrestrial backbone links. However, in this case, end-to-end delay may become unacceptable.

- A hybrid satellite/terrestrial backbone network is cost-effective and provides acceptable delay performance. In this case, the minimum cost configuration for AUTODIN II - Phase I involves installing ground stations at 6 out of the 8 proposed AUTODIN II - Phase I switch locations. This results in a savings of about 25% of the backbone line costs.

Expansion Capabilities

- Networks with an order of magnitude more traffic than AUTODIN II - Phase I can be implemented using the same basic architecture and concepts as the AUTODIN II - Phase I system and can provide the required delay, reliability and throughput. Such networks can achieve very high reliability and can be operated at lower costs per unit of traffic than smaller networks. The economics of these networks are discussed below.

- Communication costs per megabit of transmitted data in AUTODIN II - Phase I should range between 23 and 25 cents. This cost level is comparable to costs projected by NAC for ARPANET and similar networks as early as 1972.

- Economies of scale, which have previously been demonstrated within ARPANET and other medium size packet networks, exist within the large network DOD environment. A tenfold increase in traffic from the AUTOODIN II - Phase I 1.26 Mbs requirement results in about a 45% decrease in cost per transmitted megabit of data to on the order of 14 cents per megabit.

- It is the opinion of the authors that this study and past efforts demonstrated that packet switching can provide cost-effective, reliable, and responsive performance for DOD data networking. Furthermore, the benefits of packet switching, first demonstrated

for ARPANET, extend to AUTODIN II - Phase I and to networks with an order of magnitude more traffic than AUTODIN II - Phase I. Thus, approaches adopted today for current DOD requirements can be expected to survive the addition of new requirements.

DEPARTMENT OF DEFENSE IMPLICATIONS

The results of this report have long and short term implications for the Department of Defense. For the long term, the results identify the architectures and communication technologies which are cost-effective in meeting present and future DOD data communication requirements. For the short term, the algorithms and computational techniques developed by NAC can be used for the evaluation and optimization of present DOD communications systems and provide designs for implementation of new systems. Moreover, these techniques have been transferred to the Defense Communications Agency where they are being employed for this purpose. Furthermore, the studies identified potential problem areas for which capital investment will be beneficial in terms of improving the technology and minimizing development time, and the enhancement of computational techniques to enable analysis and design of very large communication networks.

IMPLICATIONS FOR FURTHER RESEARCH

It would be productive to continue further research to develop tools and tradeoffs for the study of large network problems incorporating evolving switching technologies. A major area of future research is the feasibility and cost-effectiveness of integrating voice and data applications. Tool development is also desirable to improve local access schemes such as packet radio, and for the application of domestic satellites in broadcast mode for backbone networks. The potential of these areas to the DOD establishes a high priority for these studies.

TABLE OF CONTENTS

TABLE OF CONTENTS

FIGURES

TABLE OF CONTENTS
FIGURES (Cont'd)

1. LARGE NETWORKS FOR DEFENSE COMMUNICATIONS: STRATEGIES, PERFORMANCE MEASURES, AND COST TRENDS

1.1 GENERAL

Packet-switching technology is creating the opportunity to interface, replace, and integrate many of the independent Command, Control and Communications (C^3) networks within the DOD, including AUTODIN II, SATIN IV, COINS, ARPANET, and DSCS II. Very large packet networks within DOD are a near certainty to occur. Without a thorough understanding of the cost/performance tradeoffs for such systems, DOD decision makers will find it difficult, if not impossible, to justify taking even the initial steps to determine how to best satisfy their communication network requirements.

The goals of the research summarized in this document were to develop a multidimensional picture of cost and performance tradeoffs for large scale integrated networks including satellites, ground packet radio, and terrestrial links, and to evaluate packet radio and other alternative technologies for local and regional data distribution to find conditions under which one technology is preferable to another.

The research undertaken to achieve the goals stated above built upon results obtained during past contracts with ARPA. During previous contract years, a variety of tools were developed to allow economical cost/performance tradeoff studies. These tools included: basic analysis and design algorithms for optimization of terminal processor location, topological optimization, throughput and delay analysis, and reliability analysis procedures. At the onset of the current contract period, these programs operated in stand-alone mode and thus required integration to complete the large network studies. Studies performed in past years with the stand-alone tools included: the practical demonstration that low cost terminal access can be

achieved by hardware multiplexing at ARPANET Terminal Interface Processors (TIP's); the proof of the cost-effectiveness of multipoint lines for connecting low and medium speed terminals into ARPANET; the demonstration of the use of software demultiplexing as a means of increasing the terminal handling capacity of a TIP by a factor of 10; the theoretical calculation of capacity, error rates and delay for a system incorporating broadcast packet radio techniques on a wideband coaxial cable local distribution network servicing a large suburban population; the evaluation of the impact of satellites on a 40 node ARPANET, the establishment of the feasibility of a 1,000 IMP packet-switched network using terrestrial links, and studies of the cost-effectiveness of packet-switching within an environment containing several thousand terminals. These studies established methodology and techniques for handling large numbers of terminals and processors and various packet access methods implemented within different hierarchy levels of a large C^3 network. They also laid the groundwork required to complete the evaluation of cost and performance tradeoffs in large scale integrated packet-switched networks.

The elements of a large data network which impact its cost and performance include:

1. Number of Host computers and terminals,

2. Number of geographically distinct installations,

3. Geographical spread of Host and terminal requirements, and

4. Traffic volume and characteristics.

In this study, we are concerned primarily with large distributed network which generally result from the integration of several centralized and distributed subnetworks for the purpose of achieving communications line economies and inter-network communications capabilities. Important examples of this integration trend in the Department of Defense can be found in the proposed AUTODIN II Defense Communications network, the SATIN IV System, and the PLATFORM Network. To provide specific as well as general results, a major DOD network requirement, the AUTODIN II - Phase I System - was selected (in conjunction with ARPA) for detailed development of cost and performance tradeoffs.

Results of NAC's past studies of large networks indicate that the most cost-effective structure for a large distributed network is a multilevel hierarchical structure consisting of a backbone network (high level) and local access networks (low level). The backbone network is characterized by distributed traffic requirements and is generally implemented using the packet-switching technology. The backbone network itself may be multilevel, incorporating, for example, terrestrial and satellite channels. Local access networks, on the other hand, have in general, a centralized traffic pattern (most traffic is to and from the gateway backbone node) and are therefore implemented with centralized techniques such as multiplexing, concentration, and polling. The low level local networks may also be hierarchically structured.

This volume defines the scope of the large scale network problem and summarizes some of the principal results. Discussed are:

1. Alternative large network architectures - including various local and regional architectures such as multiplexing, concentration, multidropped lines, rings, packet radio, and broadband coaxial cable systems.

2. The criteria for performance evaluation - including time delay, throughput and reliability.

3. The formulation of the large network design problem - requirements, constraints and objectives.

4. AUTODIN II Case Study requirements - including numbers of terminals and Hosts, traffic, delay and reliability.

5. Local access network tradeoffs for viable technologies as a function of critical network parameters - component cost, throughput, reliability and delay.

6. Overall network performance and cost characteristics - sensitivity of results to changes in hardware and communication line costs, use of satellite and terrestrial line alternatives.

Detailed discussions of the basis for the results reported here can be found in Volumes 2-5 of this report as well as NAC prior Semiannual Reports to ARPA. Issues not addressed in this work include security, management, and organizational issues involved in implementing fully distributed integrated networks.

1.2 STRATEGIC CHOICES FOR BACKBONE NETWORK IMPLEMENTATION

In the following two sections we describe the network alternatives and issues that must be considered. Throughout our efforts, we have assumed that the *backbone network* is operated using the packet-switching technology. It has been demonstrated in NAC's prior Semiannual Reports to ARPA that packet-switching is a reliable, and cost-effective solution for distributed data communications requirements consisting of a mix of interactive and bulk data traffic. If the traffic is predominantly file transfers and digitized speech, the best architectural strategy has not yet been determined. Investigations of this issue, and more generally, the comparison of circuit versus packet-switched disciplines, should be a topic for further study.

Within the packet-switching concept, several architectural alternatives are available and must be considered for a cost-effective backbone design.

1.2.1 Channel Alternatives

The main issue here is whether to use terrestrial links, satellite links, or a combination of terrestrial and satellite links. In the latter cases, the satellite channel access technique must also be selected from several possible alternatives (point-to-point, TDMA, FDMA, slotted ALOHA, etc.).

Terrestrial links are available under several different service offerings, e.g., Telpak, DDS (Digital Dataphone Service), Wideband Service 8000, and T1 carrier. For a given connection, the most cost-effective selection from among the available services must be made.

In general, the selection of the best communications carrier is not based on line cost only, but on channel delay, reliability, and error quality considerations. Furthermore, the use of some carriers (e.g., satellite and T1 carriers) may require the development of specialized software and hardware interface capabilities in the nodal processors. In this case, the additional nodal cost must also be considered. It is not possible to state general conclusions concerning the best channel strategy since this is critically dependent on the amount of traffic to be accommodated by the channel.

Finally, the selection of the carrier may have an impact on the selection of backbone node locations. For example, if DDS service is used, backbone nodes will perferably be installed in "digital" cities. Similarly, if domestic satellite service is used, backbone nodes will perferably be installed near carrier ground stations.

1.2.2 Node Alternatives

A conventional packet switch (i.e., ARPANET IMP) has well-defined limits in terms of throughput and the number of modem interfaces it can support. In a large network implementation, nodal throughput and interface requirements often exceed such limits and necessitate the development of new nodal architectures with upgraded capacity.

One possible means to increase node capacity consists of combining several conventional nodes in a fully interconnected cluster. The cluster can be viewed as a supernode with an external throughput and modem interface of higher capacity than the single IMP capacity. Obviously its capacity is lower than the sum of the capacities of all the individual IMP's in the cluster, since some of the interfaces are used for internal connections, and transit traffic typically traverses two IMP's in the cluster.

Another approach to higher nodal capacity is the multiprocessor switch implementation (e.g., Pluribus IMP). The multiprocessor switch is a device with several processors, I/O, and memory modules interconnected together and performing different packet switch functions in parallel. The multiprocessing environment and the overlap of different functions lead to higher throughput and modem interface capability.

Important issues in the selection of a switch architecture are the degree of centralization versus decentralization, the impact of survivability questions on the desirable number of switches, and the technological risk and cost of developing new multiprocessor switch hardware and software as opposed to using existing devices within nodal clusters.

1.2.3 Topological Structure Alternatives

Topological alternatives range from star-like centralized systems to fully distributed grid-like networks. For a small or medium size backbone network (say up to 20 or 30 nodes), the typical structure is homogeneous, with identical hardware, software, and functional requirements at nodes. For networks of a larger size, a two-level hierarchical topological structure may be more cost-effective [NAC, 1973] and should be considered as a possible alternative.

In a two-level packet network, the high level net has higher node and link capacities than the low level "subnets." A subnet is a packet network connected to the high level net through one or more high level nodes which act as gateways. Communications between different subnets is usually achieved through the high level net. However, neighbor subnets may communicate with each other directly if routing protocols permit.

The cost-effectiveness of a two-level packet network implementation is obtained by:

1. Use of high bandwidth cost-effective communications offerings (e.g., satellite channel or T1 carrier) for long distance and high capacity trunks between high level switches, and

2. Reduction of local access costs from terminals to backbone switches obtained by providing a large number of geographically well-distributed high level switches.

1.3 STRATEGIC CHOICES FOR LOCAL DISTRIBUTION

In the last section, we discussed alternatives for the backbone network. A wider range of alternatives is available for the local distribution networks. In the most general case, a given overall network may use different strategies for different portions of the overall network.

1.3.1 Host and Terminal Access

In a large integrated data network, local access from terminals and Hosts to the high level backbone network can be implemented with a variety of different techniques. The available alternatives which are considered in this study are:

- Dialup.
- Dedicated point-to-point lines.
- Polling on a multidrop circuit or on a loop.
- Time or frequency division multiplexing.
- Concentration (local or remote).
- Distributed ring structure.
- Distributed store-and-forward structure.
- Packet radio broadcast.
- Packet satellite broadcast.
- Cable television.

Of the above alternatives, only a subset is generally feasible for given user requirements (e.g., response time, reliability, security, mobility, etc.). For example, if terminals are mobile or rapidly deployable, radio or satellite broadcasts are the only feasible alternatives. A direct dedicated connection from a fixed location Host to one or more backbone nodes is the most commonly used connection strategy, but is only one of the viable options.

In practical applications, two or more colocated Hosts may be connected to a common Front End Processor (FEP) which, in turn, is connected to the backbone node via a dedicated line. This case clearly reduces to the previous one by simply merging the colocated Hosts into a Host cluster. In the design phase, the best strategy is selected from the set of feasible alternatives, based on trade-offs between cost, performance, and growth flexibility.

1.3.2 Segregation vs. Integration

An important issue is the one of using separate networks to satisfy different groups of users versus combining requirements and serving all groups via an integrated common user network. It may be observed that in some applications, there are subsets of terminals that communicate predominantly with only one Host. In this case, the most appropriate local distribution strategy may correspond to the direct connection of the terminals to the Host, rather than to a backbone node. In the limit, if the large terminal and Host population can be partitioned into several disjoint subsystems and communications requirements exist only within each subsystem with no intercommunications between sybsystems, the best strategy may be a segregated (each subsystem is connected via a private network) rather than an integrated strategy (all subsystems share a common backbone network).

Most of the tradeoff studies and design guidelines are based on the assumption that most Host-to-Host and terminal-to-Host traffic is switched through the backbone network. However, the tradeoff between segregation and integration was thoroughly investigated for the specific AUTODIN II application. For that case, it was concluded that network integration is the preferred approach.

1.4 PERFORMANCE MEASURES

There is little difference between the performance requirements for small and large networks. What is different is the degree of difficulty with which performance can be evaluated. Since the main purpose of this effort is the investigation and comparison of alternatives (rather than the control) for large network architectures, we identify here the performance measures which are directly related to network architecture and must, therefore be considered during the design phase.

Given the set of network requirements (number and location of Hosts and terminals, traffic patterns, traffic characteristics, etc.), the performance of the network will depend, in part, on the hardware selection and architectural design (e.g., line layout, buffer size, line speed, processor speed, etc.) and, in part, on the communications software design (line protocols, routing protocols, end-to-end protocols, etc.).

Some of the performance measures depend more critically on hardware selection and architectural design (e.g., path reliability) while others are more directly related to the communications software (e.g., flow control, congestion protection, stability, etc.).

1.4.1 Delay

The delay from terminal to Host in a hierarchical network structure is equal to the sum of the delays in the local access segments (terminal-to-backbone and backbone-to-Host) plus the delay over the backbone segment. Since delay is a random variable, average values or probability distributions must be considered. For example, network delay could be defined therefore as a maximum of the average delays evaluated over all node pairs. Unfortunately,

this definition leads to a cumbersome procedure for the evaluation of the delay, requiring the examination of all terminal-to-Host pairs. Furthermore, the constraint on maximum delay is difficult to account for in the topological optimization problem.

For design purposes, it is more effective to use a definition of network delay equal to the sum of the maximum times spent in the local distribution networks and the average time spent in the backbone network. For the design, delays are evaluated during peak hour traffic conditions and for the average block length of a typical transaction. Average rather than maximum delay is used for the backbone segment since the former is much easier to compute for design and evaluation purposes in a distributed packet-switched network. The approximation is quite acceptable in practical applications since backbone delays are generally much lower than local access delays and have marginal impact on global delay performance. In practice, a large network may support several different types of terminals and applications; and a given terminal may only be allowed to communicate with a limited number of Hosts. To handle this general case, we can consider a vector of delay values, where each entry corresponds to a given type of application. For example, we may have different values of delay for interactive type traffic, RJE transfers, and Host-to-Host communications.

The existence of different delay requirements for different terminal types and applications is an important variable in the design of the local access subnetwork since some access strategies may not be feasible for a given terminal type under given delay requirements. This means that particular attention must be given to this performance measure during the local access design stage. Backbone network design, on the other hand, has been found to be insensitive to the presence of different delay requirements in the local subnetworks, provided that the average delay in the backbone network is small compared to end-to-end delay requirements.

1.4.2 Throughput

Throughput performance depends primarily on network topology, traffic pattern, and delay requirements. Since traffic pattern knowledge is often not accurate in large data network applications, throughput performance must be associated with the sensitivity of network throughput to variations of traffic patterns. It is generally true that the pattern of traffic (i.e., percentage of the total traffic moving between specific pairs of nodes) can be predicted with more accuracy than the actual amount of traffic to be sent. This phenomena is true because one generally can predict such patterns by knowing the locations and numbers of users, and specific applications. Thus, when evaluating throughput, one wishes to know the maximum total traffic that can be sent when this traffic follows a specific pattern.

The throughput capability of a network is thus calculated for a specified pattern of input traffic requirements, and is defined to be equal to the amount of traffic that can be sent according to this pattern when the delay (which is an increasing value of input rates) reaches its maximum admissible value. The sum of the input rates that yields this delay is defined to be throughput for the network under a specified traffic pattern.

1.4.3 Reliability

In a network with less than perfectly reliable nodes and links, the reliability of communications between a terminal and a Host computer is defined as the probability that there exists at least one path connecting them. Network reliability may be defined as the minimum of all node pair reliabilities (i.e., the reliability between the most unfavorable node pair) or as the average over all node pairs. In reality, both minimum and average (and possibly

distribution vs. fraction of node pairs) are important to fully characterize network reliability performance. In this study, the minimum is chosen as the measure of reliability. This measure, while appropriate, is exceptionally difficult to compute for large networks. Therefore, a conservative estimate (i.e., one that is always lower) is used which is easier to evaluate. This estimate is the product of the reliabilities of backbone and local access subnets. Thus, the definition of network reliability can be extended to handle problems with different subsystems having different reliability requirements (in which case, a different value is calculated for each subsystem).

1.5 PROBLEM FORMULATION

Having specified the performance measures which serve as constraints to the network design process, we are now in a position to formally state the network design problem. The topological design problem for a two-level hierarchical network can be formulated as follows:

Given:

1. Terminal and Host locations.

2. Matrix of traffic requirements (terminal-to-Host and Host-to-Host) which can be combined to provide a throughput requirement.

3. Delay requirements (possibly different requirements for different subsystems).

4. Reliability requirements (possibly different requirements for different subsystems).

5. Candidate sites for backbone nodes.

6. An array of available component alternatives and their costs (e.g., line tariff structures, nodal processor costs, hardware costs, etc.).

Minimize*:

Total communications cost D:

D = (backbone line costs) + (backbone node costs)
+ (local access line costs) + (local access
hardware costs).

Such that:

1. Throughput requirements are met.

2. Delay requirements are met.

3. Reliability requirements are met.

The global design problem consists of two subproblems:

1. The design of the backbone network.

2. The design of the local distribution networks.

The two subproblems interact with each other through the following parameters:

1. Backbone node number and locations.

2. Terminal and Host association to backbone nodes.

3. Delay requirement for backbone and local access networks.

4. Reliability requirement for backbone and local access networks.

* In some cases, it may be desirable to minimize only a portion of the total cost. For example, if previously purchased existing hardware is to be combined into a new network, the cost of this hardware could be excluded from the optimization.

Once the above variables are specified, the two subproblems can be solved independently. This problem structure typically leads to heuristic solutions of an iterative nature.

To illustrate the nature of the design process, we give an example of an iterative approach: each subproblem is optimized after assigning reasonable values to parameters 1 through 4; the parameters are modified in the direction that yields maximum cost savings; the optimization of the subproblem is carried out under the new parameter assignment; if no further reductions in cost can be achieved, the process terminates; otherwise, we select a new set of parameters and reoptimize the subproblems.

In large network applications, the global design appears to be insensitive to the selection of backbone delay and reliability requirements. In fact, the high volume and distributed pattern of traffic demands in the backbone network generally require high backbone trunk capacity (with transmission delay substantially lower except when using satellites, than the delay on local distribution lines), high backbone connectivity and high nodal redundancy. As a consequence, the variable portions of delay and reliability in the backbone net are typically one order of magnitude better than in local distribution nets. Therefore, variations of backbone delay and reliability requirements have little effect on global performance.

The new variables that characterize large network design and distinguish it from traditional network problems are the partitioning and backbone node location. One of the objectives of NAC's study was to determine the sensitivity of the total cost and performance to variations in the number and location of backbone nodes and the partitioning of terminals and Hosts.

1.6 CASE STUDY: INTEGRATED DEFENSE COMMUNICATIONS NETWORK

The previous sections have described the alternatives and the performance measures which must be considered during the large network design process. We now illustrate these concepts via the specific AUTODIN II requirements which are the basis for our case study. The data communications requirements within DOD arise from a variety of different applications. Presently, a separate communications network is often dedicated to each application, and the system consists of the superposition of several different networks which typically cannot communicate with one another.

In order to obtain better line and hardware economies, and a more flexible internet communications capability, the DOD is planning to integrate most of the data transmission requirements on AUTODIN II, a new data network implementing the packet-switching technology. Among the important decisions required during the network planning process are:

1. Number, location and structure of the backbone nodes.

2. Carrier offerings to be used for the implementation of backbone trunks (terrestrial and/or satellite).

3. Local access techniques.

AUTODIN II is a representative example of large distributed data networks, which is used as a case study in the present investigations. Therefore, design criteria and cost-performance tradeoffs developed on general network models have been applied to AUTODIN II requirements. An extensive analysis of many design alternatives for this system is given in [NAC, 1975]. The results

developed there are used as a starting point for the studies described in Volumes 3 and 4 of this report, and summarized in this volume. For easy reference, the Executive Summary to [NAC, 1975] is reproduced in the Appendix.

A summary of Defense Communications requirements and characteristics considered in the study is presented below:

A. AUTODIN II will consist initially of 35 systems, of which 3 are presently implemented with distributed computer-to-computer networks (AUTODIN I, WWMCCS and ALS) and 32 are implemented with centralized, terminal oriented networks.

B. The 35 systems include 87 Host computers and over 1,000 terminals, nationwide distributed and installed at over 250 distinct locations.

C. Total net input traffic (sum of all terminal and Host requirements) is estimated to be 1,250 Kbps.

D. Delay requirements vary according to the application. In this study, delay requirements are assumed to be less than or equal to one second for Host-to-Host communications, and less than or equal to two seconds for terminal-to-Host communications.

E. Reliability requirements vary according to the system. Systems are classified into two categories: critical systems, that must be up more than 99.95 percent of the time and non-critical systems, that must be up more than 99 percent of the time. In [NAC, 1975] we address each subnetwork reliability requirement individually. In the current effort, we require only that the 99 percent criteria be met. This allows easier abstraction of the results to other network scenarios.

F. Security issues are examined in [NAC, 1975], but are not discussed here.

For the Defense Communications requirements, the following studies were performed:

A. Cost and performance evaluation of various segregated network strategies which result when separate, private networks are used for each system.

B. Cost and performance evaluation of various integrated strategies, via common user networks, using different approaches for backbone network design.

C. Sensitivity of global cost and performance to changes in backbone node number and location.

D. Cost-effectiveness of satellite links for backbone and local access communications.

E. Integration of findings in A, B, C and D to determine multidimensional tradeoffs for the Defense Communications requirements.

In the following subsections, we summarize specific cost tradeoffs identified for the AUTODIN II example. We also consider a substantially larger test problem in order to test the extendability of some of the key conclusions of the case study. Among the items addressed are local and backbone network costs. In the local distribution area, we consider star type configurations, multiplexing and concentration, multidrop trees and loops, packet radio technology, ring technology, and the use of packet radio techniques on

CATV local distribution systems. For the backbone networks, we consider the tradeoffs between switch number and cost, the use of terrestrial backbone lines, and the use of satellites. Moreover, the sensitivity of network costs to changes in the local access, backbone lines and switch component costs are displayed.

We examine, in Section 1.6, alternatives for local access. These alternatives offer substantial room for future DOD cost reductions. It will be seen that several alternatives (packet radio, rings, and CATV) provide such opportunities. These opportunities are especially important since it will be seen (Section 1.6.2) that local access costs are on the order of 50% or more of total network cost.

1.6.1 Local and Regional Distribution Network Strategies and Tradeoffs

1.6.1.1 Stars, Loops, and Multidrop Lines

A preliminary study for the AUTODIN II System [NAC, 1975] compared two local access strategies, namely:

1. Star Connection Strategy: Point-to-point connection between each user site and the nearest backbone switch. Colocated terminals may be clustered, when cost-effective, using TDMX or concentration devices.

2. Intermediate TDMX or Concentration Strategy: This is a refinement of the star strategy, and consists of connecting user sites to intermediate TDMX devices or concentrators rather than directly to the switch, when this results in line cost savings. TDMX and concentration devices are optimally located to maximize savings.

The local access design must satisfy a minimum availability requirement of 99% up time for end-to-end connections. In order to meet such requirements, the intermediate TDMX and concentration devices of Strategy 2 must be dual-homed onto backbone switches. For the Star Strategy 1, on the other hand, no line backup is required for reliability except for those critical systems which required dual-homing under every network architecture.

Because of the additional cost for reliability enhancements, Strategy 2 is no more cost-effective than Strategy 1. This conclusion is the starting point for a family of investigations which explore the potential benefits of combining the point-to-point star scheme with a polling scheme, where low traffic terminals are polled from the backbone switches on multidrop circuits. Polling was not considered in the preliminary study because of the difficulty of implementing security in a polled environment. While the polling strategy may not be feasible for the Defense Communications System because of security considerations, we evaluate the savings that such a strategy can provide based solely on cost and performance considerations. As will be seen, polling does not produce large enough savings to warrant special efforts to overcome the security issues which currently restrict its use in the DOD.

The approach taken is to obtain a lower bound on the cost of an integrated star and polling strategy to assess the maximum savings that can be obtained. Two alternative configurations for the multidrop circuit are considered: the tree structure, and the loop structure. While the former configuration provides better cost savings, the latter is more reliable.

Figure 1.1 summarizes the results of our investigation of stars, multidrop circuits, composed of trees or closed circuits called loops. The lowest cost option is the star + multidrop tree configuration. However, the tree structure is not sufficiently reliable to meet the 99% connectivity constraint, and thus requires dialup

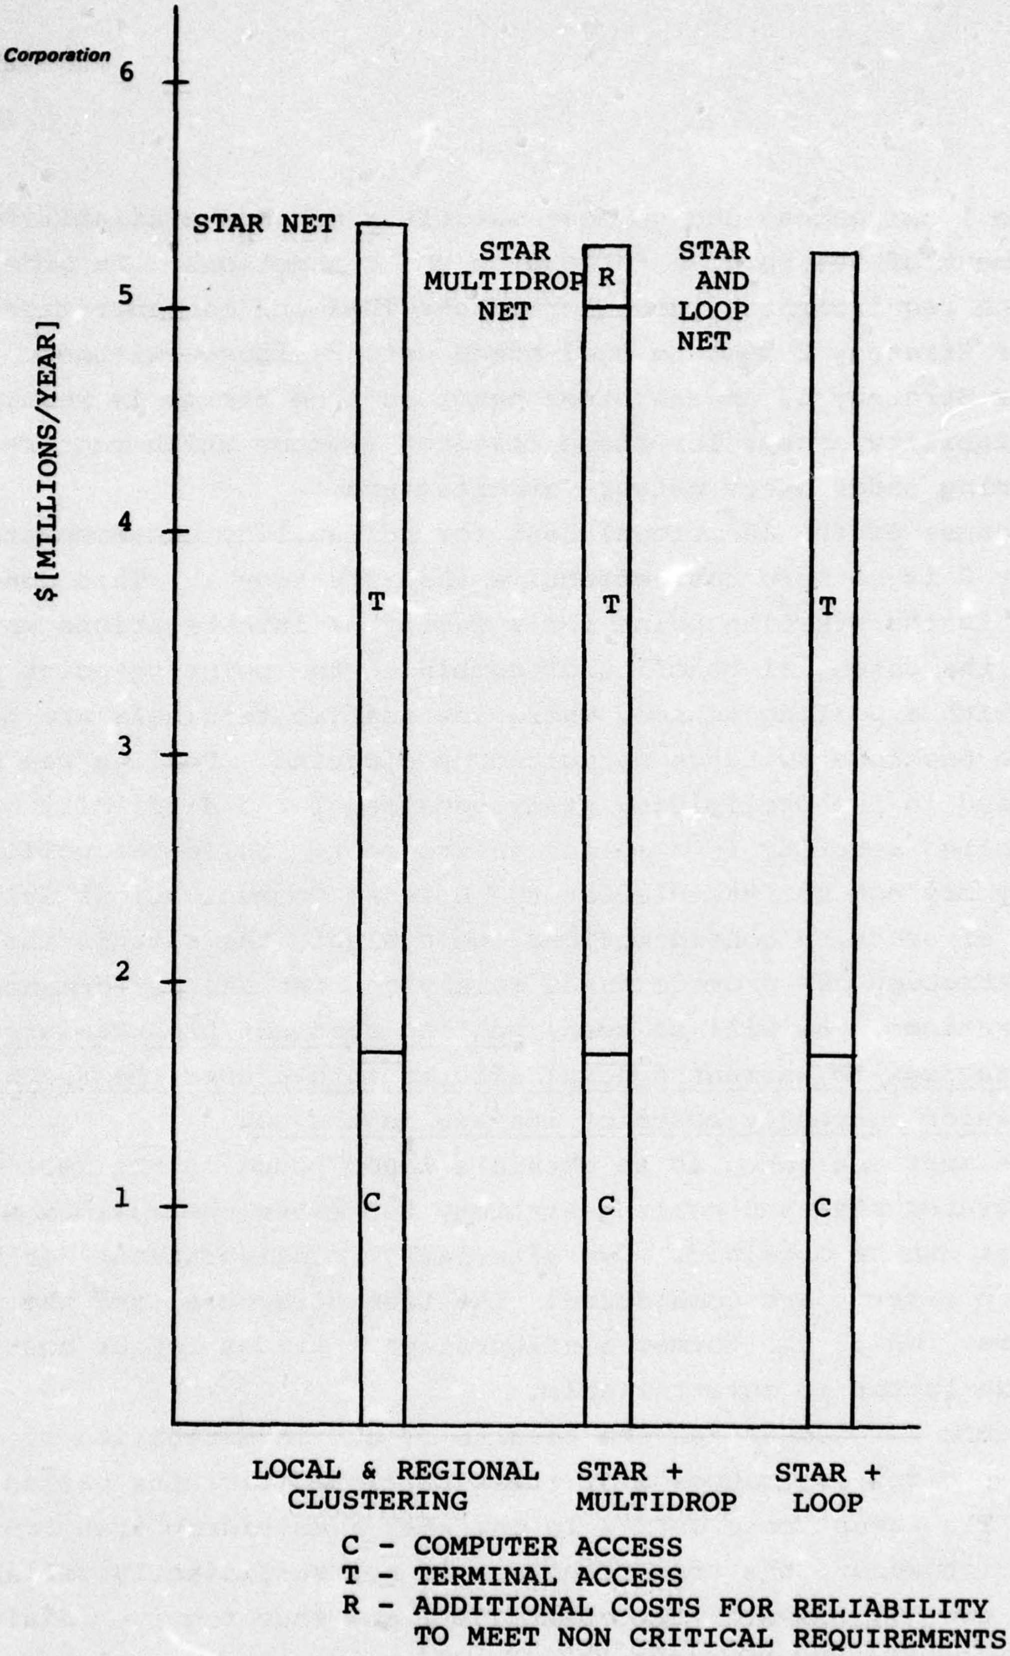

FIGURE 1.1: COMPARISON OF THREE LOCAL DISTRIBUTION ALTERNATIVES FOR AUTODIN II

backup (an undesirable feature in a Defense Network) to meet the Defense Communications standards. Adding an estimate of the least expensive method of providing dialup backup at terminals increases the cost above the next to lowest cost architecture, the loop and the tree strategy. Thus, the loop is (costwise) the most attractive among the polling alternatives. However, the cost differential between the hybrid star and loop configuration and the basic star configuration is less than 5% of the total local access cost. Moreover, this figure is an upper bound on the savings since several considerations relating to loop systems have not been included.

For a more accurate appraisal of the merits of the strategy, the following additional cost factors must be included:

1. Cost of polling software implementation in the switch.

2. Cost of providing security.

3. Cost of buffers and selectors to allow a standard polling procedure.

If we now tradeoff the marginal cost savings of 5% and the elements of potential technical risk and unknown additional costs we conclude that polling (either multidrop circuits or loops) is not a cost-effective local access technique for the present Defense Communications environment. In the future, the polling alternative may become a viable alternative if some of the following trends develop in the Defense requirements:

1. Installation of a large number of homogeneous (compatible) terminals with incorporated polling hardware and logic.

2. A large number of terminals are isolated and thus cannot be effectively time division multiplexed at a local collection point.

3. A large number of terminals have no security requirements and it is considered a safe risk for these to share a common local communication line.

Figures 1.2 and 1.3 illustrate the local access networks developed under the multidrop tree and loop alternatives.

1.6.1.2 Packet Radio For Local Distribution

The Packet Radio (PR) technology is being developed to address communication requirements of applications for which the present conventional technology is not particularly suitable [KAHN, 1975], and to provide more efficient techniques for frequency spectrum utilization. Because of conservative assumptions about the operation of packet radio systems and the primitive state of design techniques for packet radio systems when compared to the conventional communication alternatives, the analyses to be discussed are biased towards and should favor the conventional approach. Nonetheless, it will be seen that the packet radio technology can compete with conventional technologies on an economic basis for many conventional applications such as local access in AUTODIN II.

To investigate the economic feasibility of using Packet Radio in conventional military applications, we apply the technology to a set of computers and terminals of the AUTODIN II network. The objective is to assess the applicability of the PR technology and obtain a preliminary cost comparison between a conventional communication technology and the PR technology. The PR technology

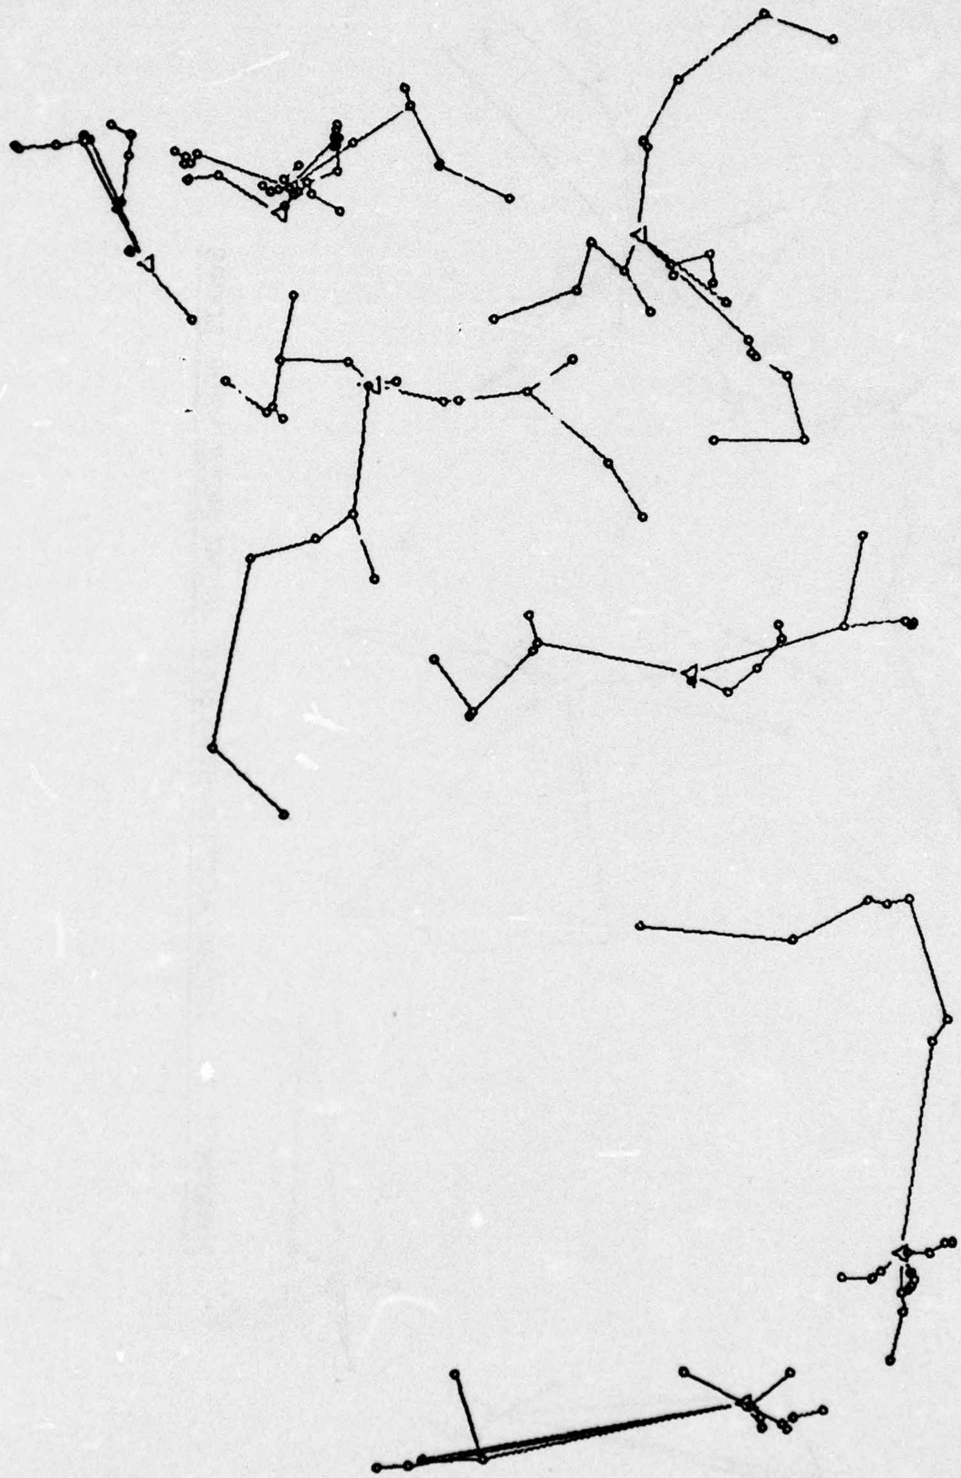

FIGURE 1.2: LOCAL ACCESS STRATEGY USING MULTIDROP TREES

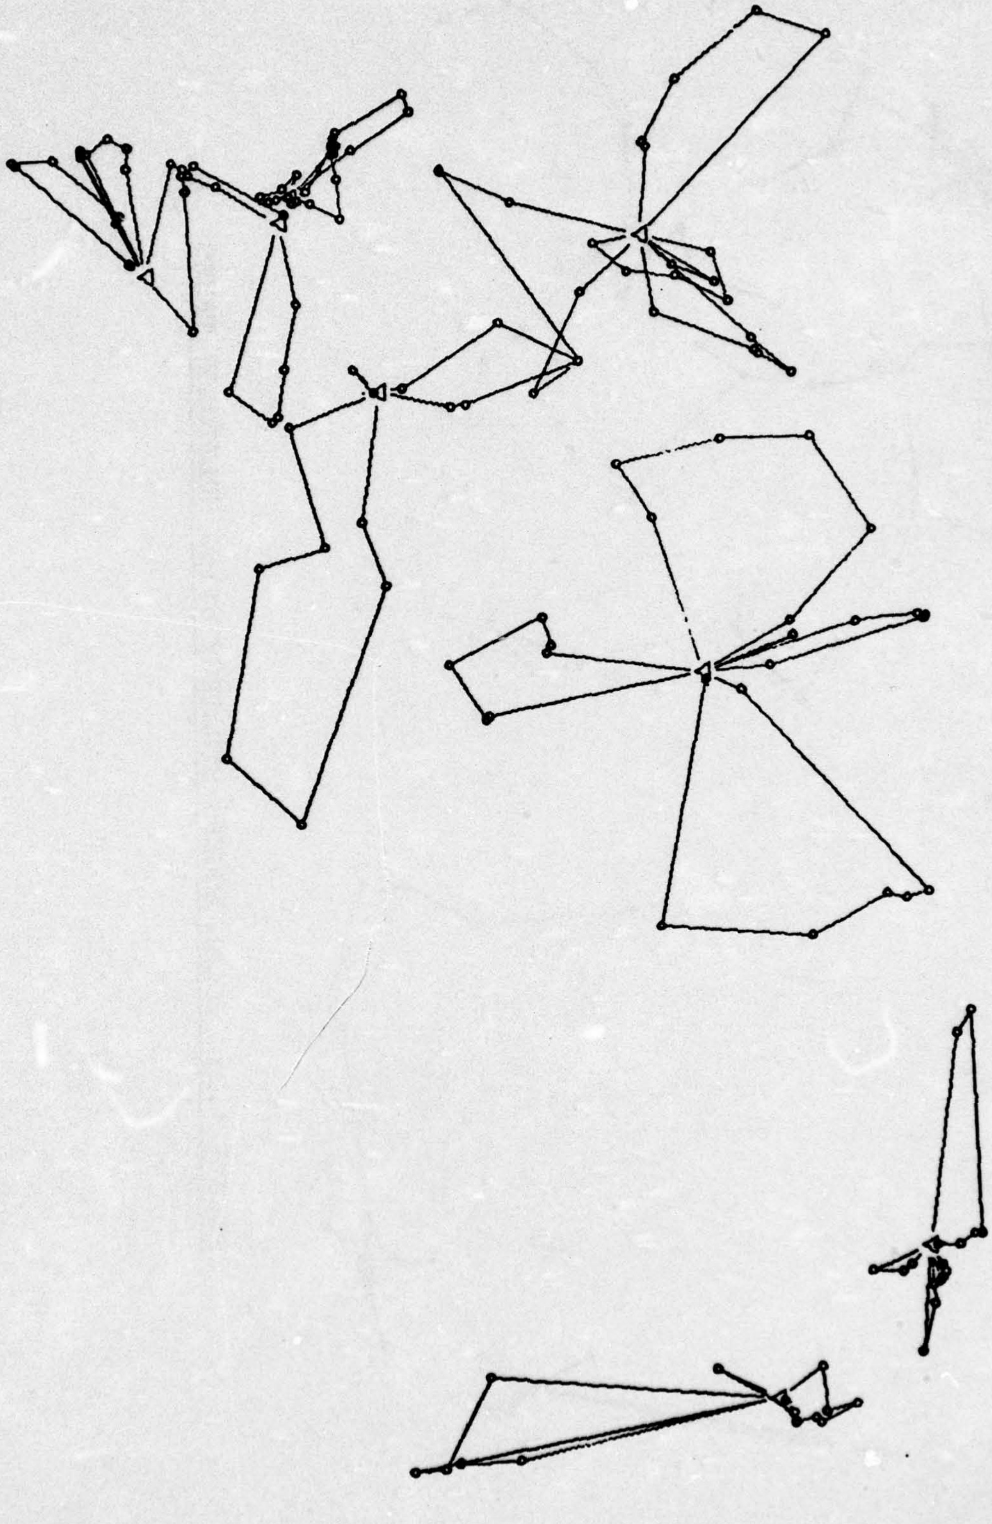

FIGURE 1.3: LOCAL ACCESS STRATEGY USING MULTIDROP LOOPS

consists of interfacing radio devices to computers and terminals and using a broadcast radio channel for connecting the set of computers and terminals. The conventional technology used as a basis for comparison consists of concentrators and multiplexers for connecting low speed terminals and computers to a backbone network in a point-to-point configuration, as discussed in the previous section.

The Washington, D. C. area was used for the case study. The area covers a rectangle of approximately 80 by 100 miles and contains a set of 20 Host computers and 109 terminals.

This case study is preliminary in several aspects. First, the PR technology has not yet been extensively tested, and its devices have not been produced on a commerical basis. Therefore, the performance of the Packet Radio Network (PRNET) is based on the modeling efforts discussed in NAC's previous reports. Since the devices have not been produced on a large scale commercial basis, the cost of communication devices is parameterized; the objective being to calculate the associated costs which make the packet radio approach equal or better than the conventional approach. Second, no reflection is provided of the "cost" of radio spectrum (a commodity which is allocated by the FCC in the public interest). Finally, computer programs have not yet been developed for optimizing PRNET design; hence, there is no claim of optimality or minimum cost design. Consequently, we are comparing unoptimized packet radio designs with optimized conventional designs - a comparison which biases the results towards the conventional approach.

There are several differences between the use of packet radio for AUTODIN II local access and the applications for which the PR technology is being developed. Terminals and Hosts are not mobile, terminals and/or Hosts are colocated, the PR Unit (PRU) or its antenna can be placed on elevated areas, and PRU's have access to (practically) unlimited electrical power sources. These differences mean that:

1. There is no need to interface each terminal or Host computer with its own PRU. A PRU can serve as a front-end to several colocated Hosts and/or terminals.

2. A PRU can perform two functions simultaneously; operate as a front-end to a set of terminals and at the same time perform the relay functions of a repeater, providing its throughput capacity is sufficient.

3. Freedom to place the PRU and the availability of electrical power enables the PRU to have a larger effective transmission range than in general deployment conditions.

Other variations which can reduce cost and improve performance are also available. For example, directional antennas at stations may significantly increase PRNET capacity. Also, a routing algorithm which enables direct PRU to PRU packet transportation without the need to go through the Station has been proposed. This algorithm is particularly useful for applications in which devices are not mobile and in which some of the origination and destination process pairs are on the radio network. The above and other available results are not utilized in this preliminary case study.

The 20 Host computers and 109 terminals in the 8000 square miles area for which the PRNET is designed are located in 25 different locations as shown in Figure 1.4. Two of these locations contain backbone switching nodes which interface to a higher level network.

For comparison of costs, cost factors are based on preliminary DCA procurement estimates for tariffed communication lines and hardware. Hardware cost factors include purchase price, installation cost (20% of base price), initial support costs (67% of base price), operations and maintenance costs (47% of base cost over 10 years), and amortization over a 10 year period with 10% yearly interest.

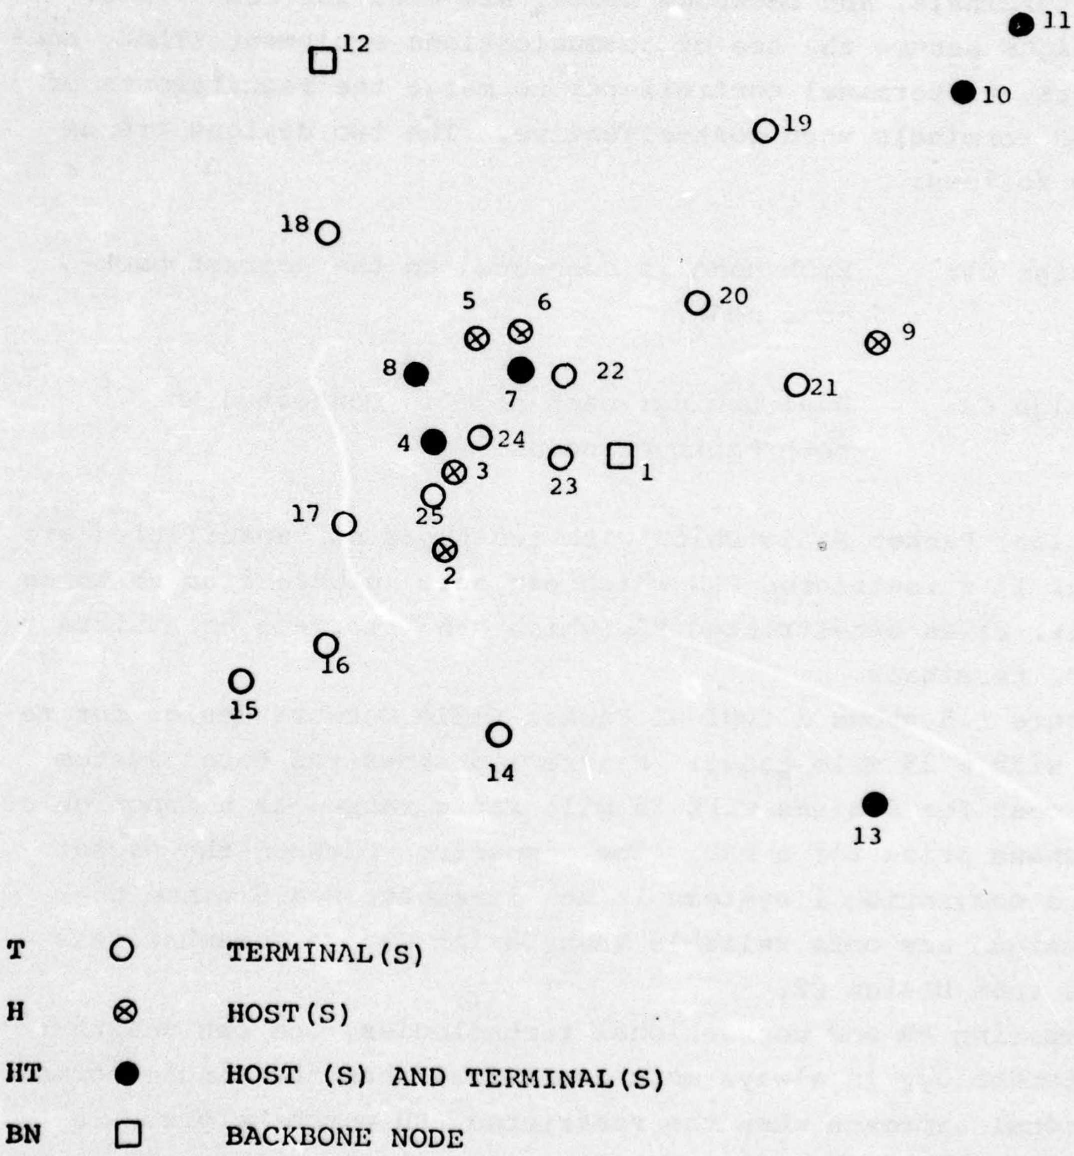

FIGURE 1.4: LOCATION OF NODES IN CASE STUDY
(WASHINGTON D.C., METROPOLITAN AREA 80x100 MILES

The cost of the PRNET designs are derived as a function of the unit price cost of a PRU. Station hardware cost is decomposed into the cost of the minicomputer PDP11/45 with an ELF operating system and 64K of core, plus the cost of a PRU and its interface.

Two conventional minimal cost designs for the same set of computers, terminals, and backbone nodes, are used for comparison. Such designs assume the use of communications equipment (TDMX, concentrators, or terminal controllers) to merge the requirements of colocated terminals when cost-effective. The two designs are defined as follows:

Design C1: Each node is connected to the nearest backbone node.

Design C2: Dual homing; each node is connected to both backbone nodes.

In addition, Packet Radio Units with two types of capabilities are examined: 1) a restricted PRU which can only interface up to three terminals; 2) an unrestricted PRU which can interface an arbitrary number of terminals.

Figure 1.5 shows a typical Packet Radio Network design for repeaters with a 25 mile range. Figure 1.6 shows the total system monthly cost for designs with 25 mile radio ranges as a function of the purchase price for a PRU. The comparison between the packet radio and conventional systems is not straightforward since the PRNET designs are more reliable than Design C1 but somewhat less reliable than Design C2.

Comparing PR and conventional technologies, one can see that the PR technology is always more economical than the single homed conventional approach when the restricted PRU purchase price is less than about $7,000 or the unrestricted PRU has a cost less than

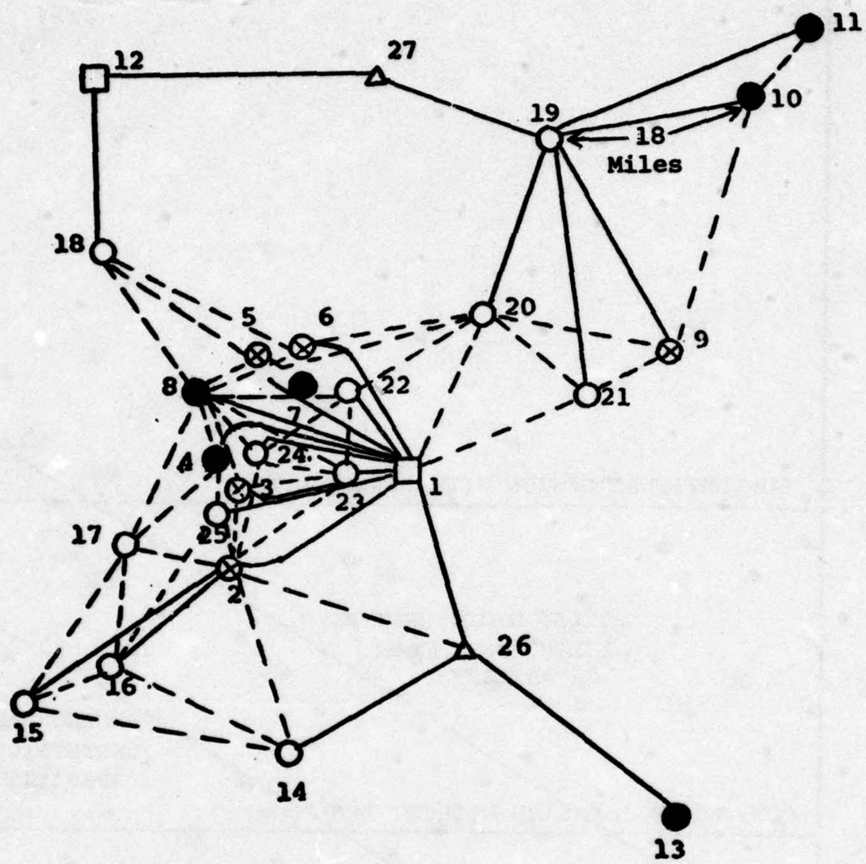

T ○ TERMINAL(S)

H ⊗ HOST(S)

HT ● HOST (S) AND TERMINAL(S)

BN □ BACKBONE NODE

PR △ PACKET RADIO UNIT (REPEATER)

———— RADIO LINKS FOR ROUTING

-------- OTHER RADIO LINKS

FIGURE 1.5: PRNET CONNECTIVITY AND ROUTING FOR PRU TRANSMISSION RANGE OF 25 MILES

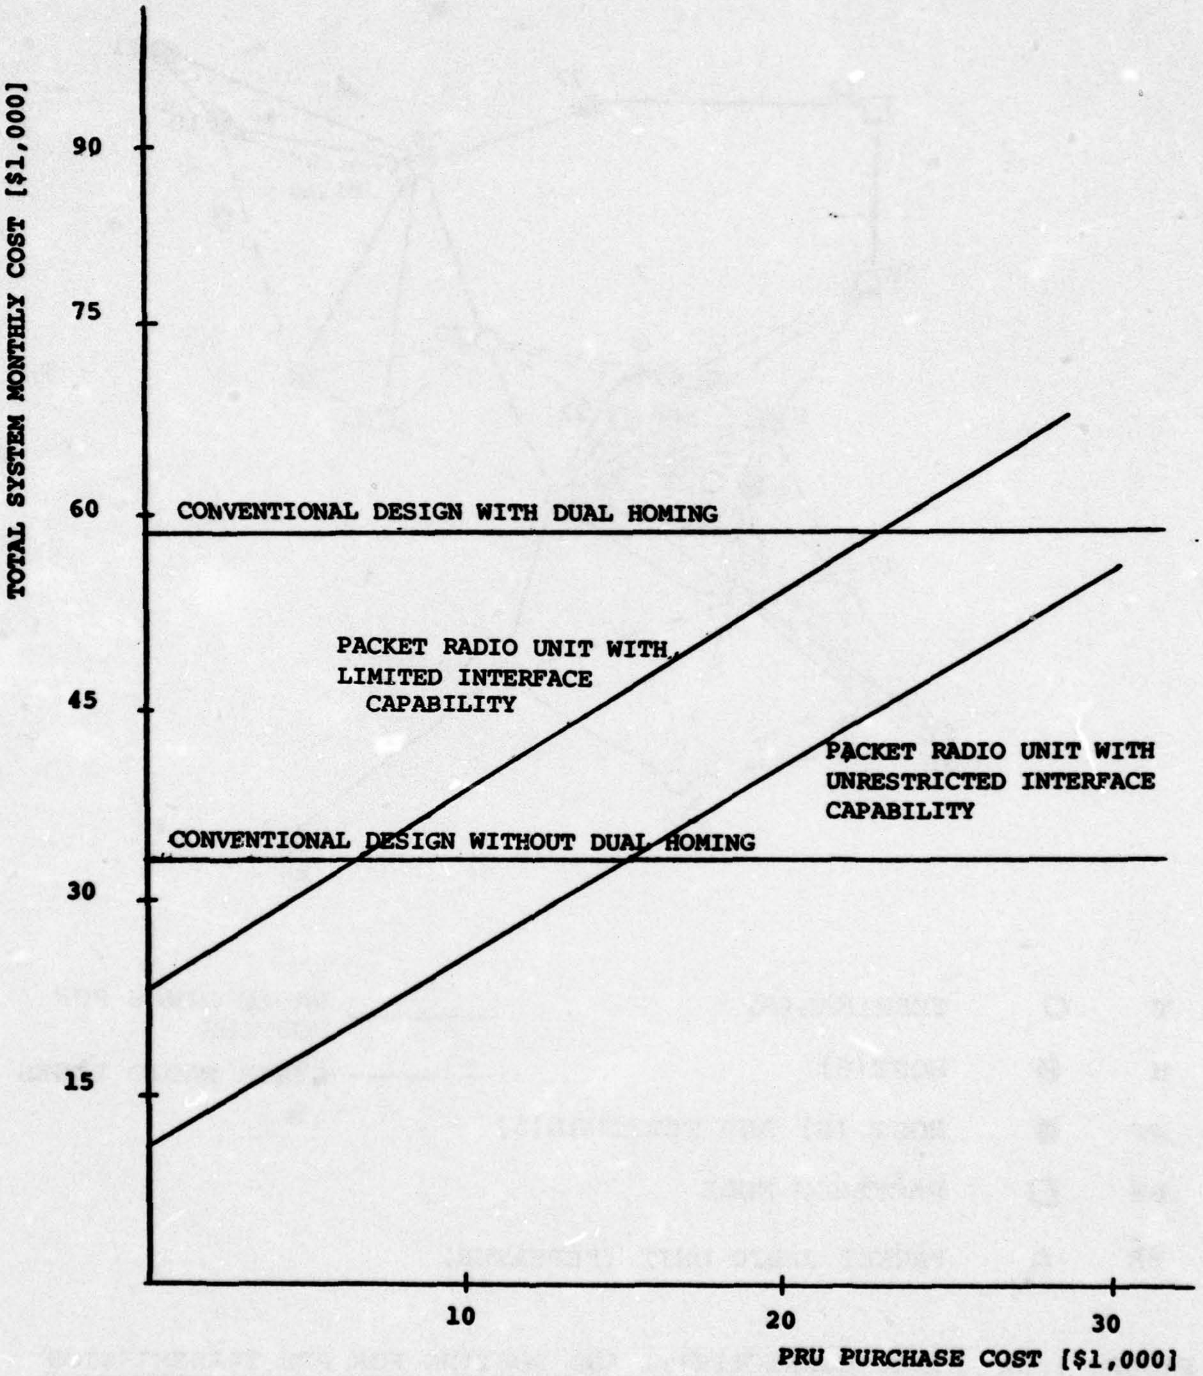

FIGURE 1.6: MONTHLY SYSTEM COST FOR ALTERNATIVE DESIGNS

about $15,000. For a restricted PRU cost of less than about $25,000, the cost of all PRNET designs are less than the dual homed design C2 while an unrestricted PRU cost less than $33,000 leads to a more economical packet radio system. These costs are within the range of current packet radio technology.

To increase the reliability of the PRNET designs to that of Design C2 requires adding several PRU's for radio connectivity, increase of backbone node capacities by using directional antennas, or adding another station and a point-to-point channel. Appropriate additions lead to increased packet radio network costs but allow a direct comparison between the packet radio and the dual homed network.

Figure 1.7 shows a comparison of Design C2 with the upgraded PRNET designs. The PRNET designs are for the two different assumptions of the number of ports per PRU and a transmission range of 25 miles. One can see that the restricted port PRU design is more economical than the conventional design for a PRU purchase price of about $20,000 or less and the unrestricted system for a PRU purchase price of $30,000 or less.

1.6.1.3 CATV For Local Distribution

We now examine the economics of cable television systems for data communications within the Washington area test case. Since, CATV (cable television) and MATV (master antenna) systems are intended primarily for the transmission of entertainment video, their characteristics are not designed for data transmission. However, with signal-to-noise ratios of better than 40 dB and bandwidth, in an unused TV channel of 6 MHz, they would seem to be ideal for data transmission. We will see that CATV systems can be very cost-effective for high density areas.

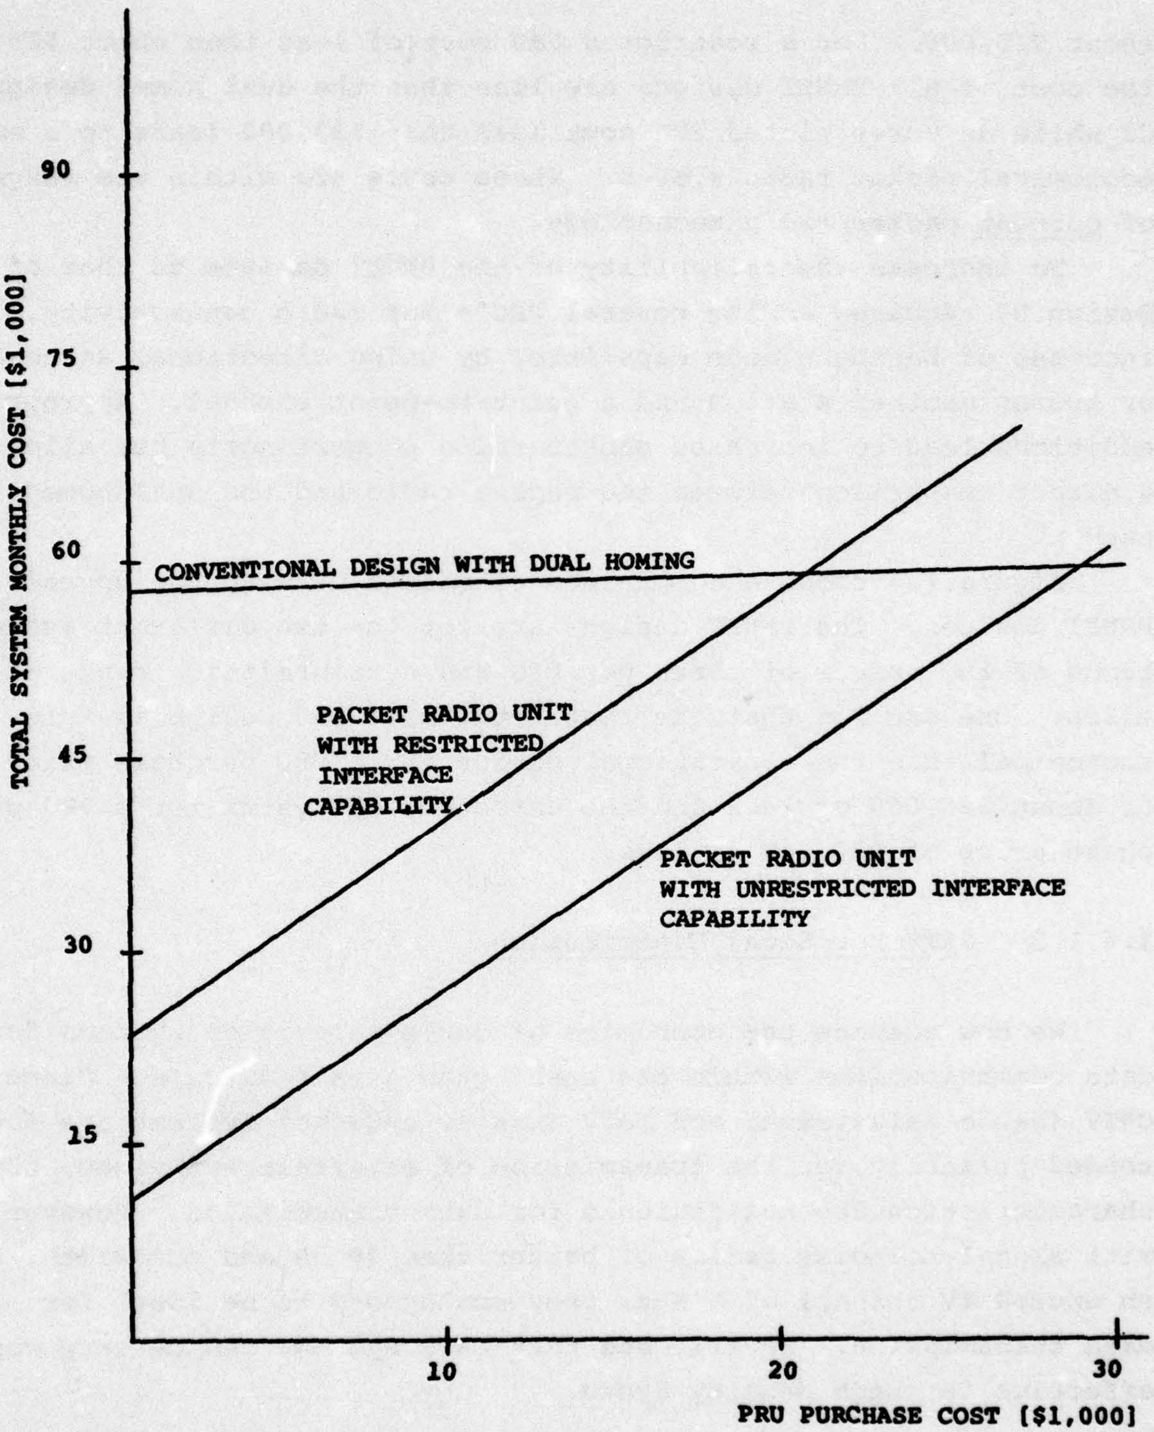

FIGURE 1.7: MONTHLY SYSTEM COST FOR ALTERNATIVE DESIGNS WITH DUAL HOMING RELIABILITY

Two-way CATV systems permit input from virtually any location in the network. The result is a large number of noise sources fed upstream toward a common "head end" on a tree network of coaxial cable and two way "upstream" and "downstream" amplifiers operating at different frequencies. CATV amplifiers have a noise figure of about 10 db for a 6 MHz channel. Cascading amplifiers can increase the effective system noise figure by 30 db or more. Nevertheless, system specifications on the signal-to-noise ratio for CATV systems are stringent enough so that data can be sent with existing analog repeaters, and no digital repeaters, such that bit rate error probabilities are negligible.

The feasibility and economics of CATV systems for data transmission is demonstrated in Volume 4 of this report. The technology is now applied to the Washington area case study and compared with conventional technology designs. The conventional designs are those reported in the previous section.

The maximum radius of a single CATV system or "hub" is about 20 miles, hence only small parts of the Washington area can be covered by CATV systems. Of the 25 nodes, 5 hubs, in typical CATV urban settings could cover 14 nodes. For each of these nodes, a terminal interface and modem is required. The cost of this unit would be $3,200 in lots of 200 and could be brought down to $500 in lots of 10,000. A head end minicomputer costing $14,000 is also required. Finally, since the hubs of the CATV system are disconnected, we must use another technology to complete the network. We can assume a conventional star connected network. A CATV system is a tree structure with no redundancy. If we were to add a redundant path from a node beyond the first power supply, the average cost in CATV equipment would be about $3,000. To provide full redundancy back to the head end there would be an average cost of about $24,000. With any redundancy a digital switch costing $2,800 is required.

Figure 1.8 shows the total monthly system costs of the CATV and conventional technology designs. For the CATV system, three designs are plotted. One corresponds to a CATV system with no added capability for redundancy. A second design corresponds to the case in which redundancy is provided so that there is an alternate path from each Host or terminal beyond the nearest power supply. Finally, the third design corresponds to the case of full redundancy in which an alternate path is provided from the Host or terminal back to the minicomputer at the head end.

For the CATV systems, the costs included are those of the digital equipment, the cost of CATV equipment to provide redundancy and the cost of the point-to-point lines to connect various CATV hubs. Not included is the rental fee charged by the commercial cable television company for use of two 6 MHz CATV channels.

In the figure all designs correspond to dual homing except the conventional design C1 and one of the CATV designs. The expected cost of a CATV modem-interface is less than $3,000. Hence, the CATV design with no redundancy is competitive with the single homing conventional design and the CATV design with partial redundancy is about $15,000 per month less than the conventional design with dual homing. The conventional design with dual homing is very close in price to the cable television design with full redundancy.

CATV systems are also cost-effective for very high density areas; in particular, the redundancy could be provided at a lower cost per terminal since a large number of terminals could share redundant paths. CATV systems are ideal for urban areas since the coaxial cable provides excellent shielding against urban noise.

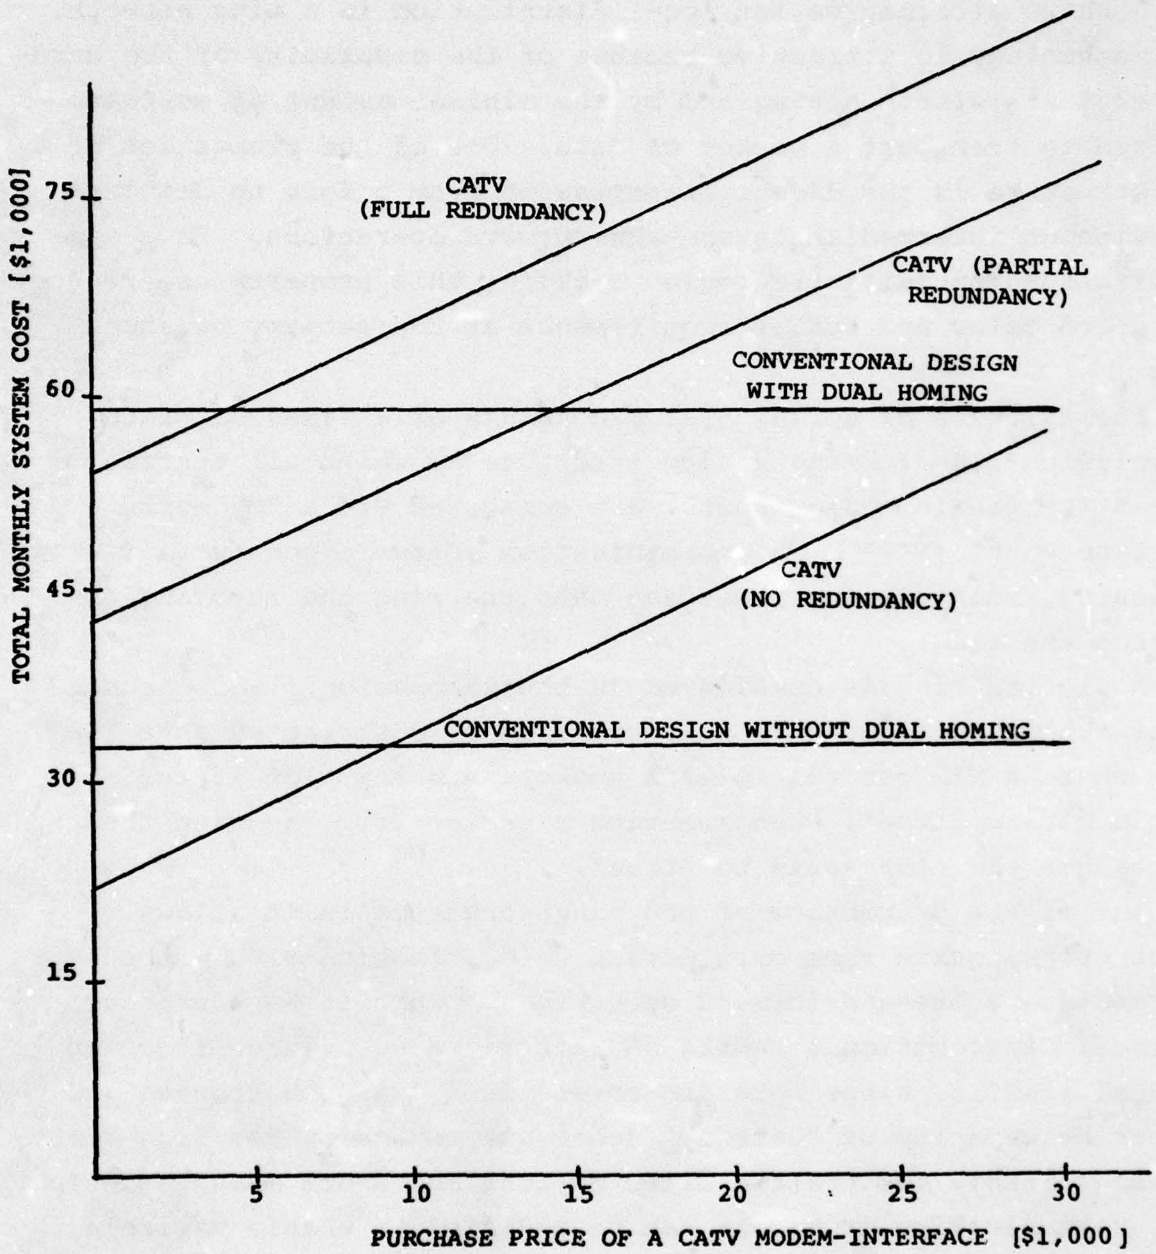

FIGURE 1.8: TOTAL MONTHLY SYSTEM COST FOR CATV AND CONVENTIONAL DESIGNS

1.6.1.4 Ring Technology for Local Distribution

A third alternative for local distribution is a ring network. This technology is attractive because of the simplicity of the hardware elements within a ring and by the minimal amount of software required to transport a packet of data. One of the properties of a ring structure is the direct transmission from origin to destination without intermediate store-and-forward operations. In a system with substantial intra region traffic, this property can reduce end-to-end delay and buffer requirements at the gateway backbone nodes.

The hardware of a ring system consists of a fixed bit rate transmission line forming a ring structure to which all traffic sources (terminals and/or Hosts) are connected via a "Network Interface Port" (NIP). The communication software consists of addressing, multiplexing a message onto the ring and removing one from the ring.

A slotted ring is considered in the discussion. The channel access scheme considered is characterized by distributed intelligence where a NIP can multiplex a packet into any slot if the slot is marked "free"; when removing a packet from the ring the NIP changes the slot state to "free".

One of the properties of the ring structure is it allows direct transmission from origination to destination NIP, without intermediate store-and-forward operations. This is an advantage for local distribution networks characterized by a large fraction of local traffic, since both end-to-end delay and the storage and processing capacity of "Gateway" nodes are reduced. The ring structure is suitable for traffic mixes of long and short messages. In fact, communication protocols can be modified to enable variable message size transmission over the ring.

For a given network cost, the throughput-delay capabilities compare favorably with other technologies [NAC, 1974A]. However, there may be a large variance between the average delays experienced by different NIP's, depending on the traffic rates and pattern. For example, let 1, 2, 3, ... be the order of NIP's on the ring corresponding to the direction of bit flow, and assume that NIP 1 has a high input rate. In this case NIP 1 will capture many of the empty slots, resulting in a high ring utilization as "seen" by NIP 2, and consequently a high average delay. The "ring domination" problem can be resolved by proper controls, which will result in a lower maximum throughput. The limitations of the ring technology are its unreliability and the inflexibility in routing.

We now examine the cost of a ring structure connecting the 20 Hosts and 109 terminals in the Washington area example shown in Figure 1.4. It is assumed that all colocated terminals and/or Hosts can interface to a single NIP. Traffic considerations show that a single ring of a 230 Kb/s is sufficient to provide end-to-end delays equivalent to those achieved with other technologies. Two parallel rings, however, are provided in the design to obtain reliability equivalent to that of "dual homing" in the conventional technology case.

The only unknown cost of this technology is that of the NIP. Hence, this cost is parametrized. Figure 1.9 shows the total cost of the ring design as a function of NIP purchase price. As reference points, the costs of the single and dual homed conventional designs discussed previously are indicated by the horizontal lines marked C1 and C2, respectively. As is evident, the ring system is always more costly than the simple star designs, but is more cost-effective than the dual homed systems for NIP costs less than about $17,000.

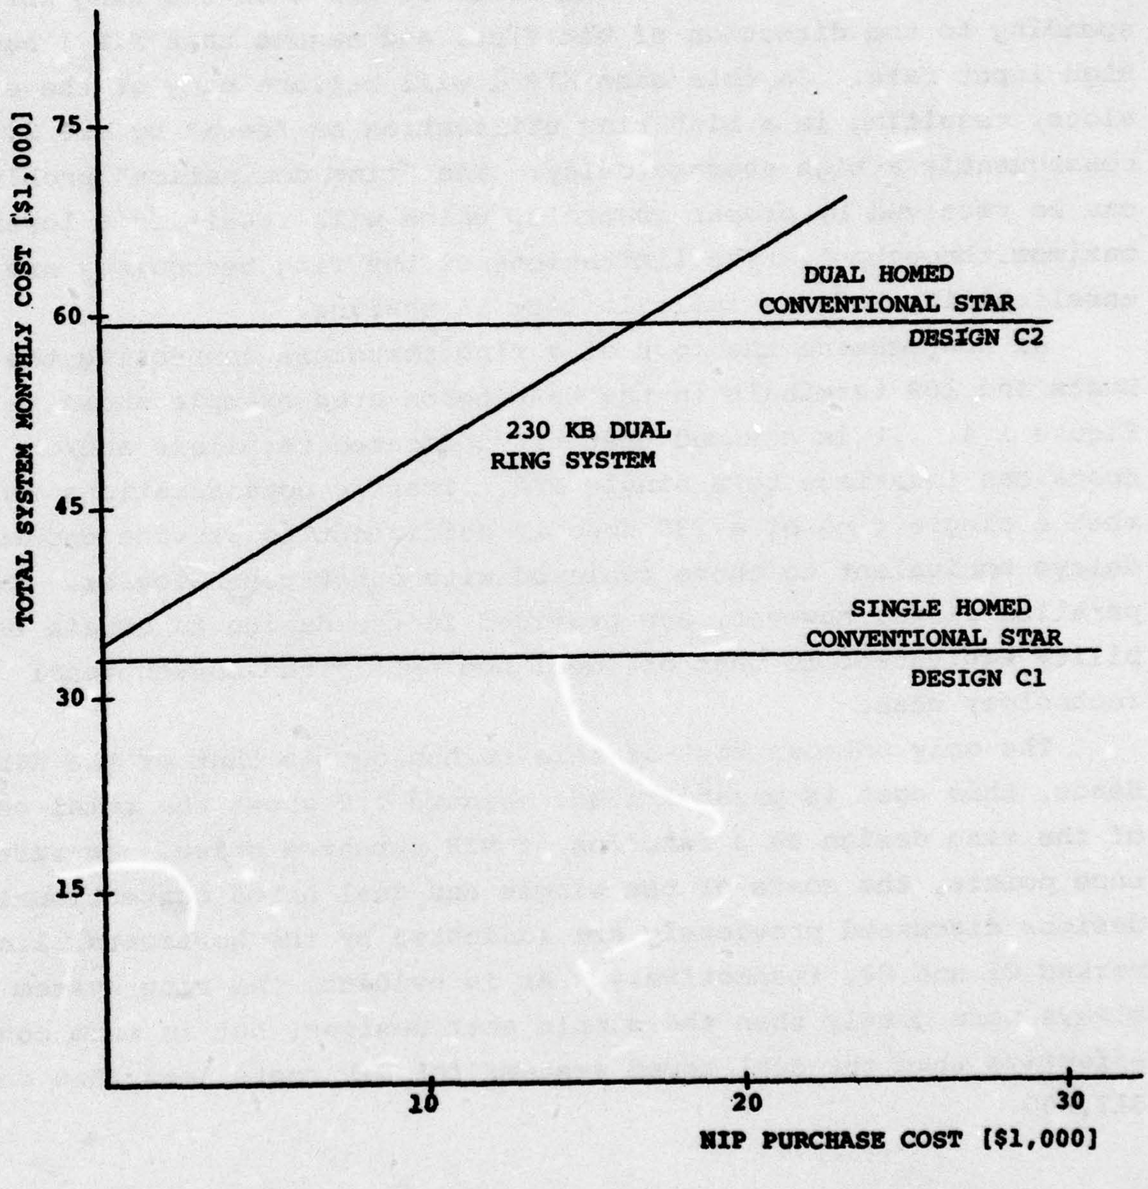

FIGURE 1.9: TOTAL SYSTEM MONTHLY COST VS. NIP PURCHASE PRICE FOR A RING SYSTEM

The assessment of the cost components for the ring design [FRANK, 1976] show that the number of nodes to be connected is the major factor determining the total cost. Hence, the technology is cost-effective in cases which may require high throughput, high reliability, and large distances between nodes, providing the number of nodes to be connected is small.

1.6.2 Cost Trends of the Total Network

1.6.2.1 System Models

In the previous sections, we have examined several alternatives for local access. We now examine overall network cost and performance trends. In order not to distort the tradeoffs, we assume that conventional techniques are used within the local access area. To investigate backbone network cost trends, a backbone switch location procedure was developed and applied to the AUTODIN II problem as well as to a substantially larger problem, called the "Growth System Model" which could potentially model the future AUTODIN II System.

The Growth System Model was developed specifically to study the impact of a tenfold growth of AUTODIN II - Phase I traffic requirements on the switch location strategy, and cost and performance tradeoffs. Terminal locations and traffic volume were generated in proportion to population distribution of the United States. The Growth System contains on the order of 450 locations generating traffic requirements of 12 megabits/second. The local access network strategies are identical for the two systems. More powerful switches and higher speed channels are required, however, for the Growth Backbone Network, since the traffic volume of the Growth System is approximately 10 times larger than that of the AUTODIN II - Phase I system. A feasible approach to building higher capacity switching facilities is to cluster lower capacity nodes within

a single switching center. Thus, the capabilities of a Pluribus IMP can be achieved by an interconnected cluster of conventional ARPANET IMP's and an ultra high throughput IMP could be built, in theory, by clustering Pluribus IMP's.

1.6.2.2 Trends for Terrestrial Network Designs

The most important trends and properties discovered during our experiments are reported below:

A. The curves of total cost versus number of backbone nodes are very flat. In the range considered here (5-20 switching nodes), the difference between the minimum point and the maximum point usually amounts to only 10 percent of the total cost. These results suggest that other military constraints such as switch security, logistics, and survivability will not prevent the building of efficient networks with total cost within a few percent of the optimum. Similarly, previous studies with highly distributed structures and large numbers of switch locations have comparable costs and performance.

B. The curves were lowest in the region of 9-12 nodes for both AUTODIN II - Phase I and the Growth Models (see Figure 1.10). Since the total traffic in the Growth System is 10 times higher than in the AUTODIN II - Phase I system, the capacity of the switches and lines must also be much higher. If the same switching and line facilities were employed in both, we would expect the optimum number of backbone nodes for the Growth System to be much higher than that for AUTODIN II - Phase I System.

Sample topologies for AUTODIN II - Phase I and Growth System with different location strategies are shown in Figures 1.11 and 1.12.

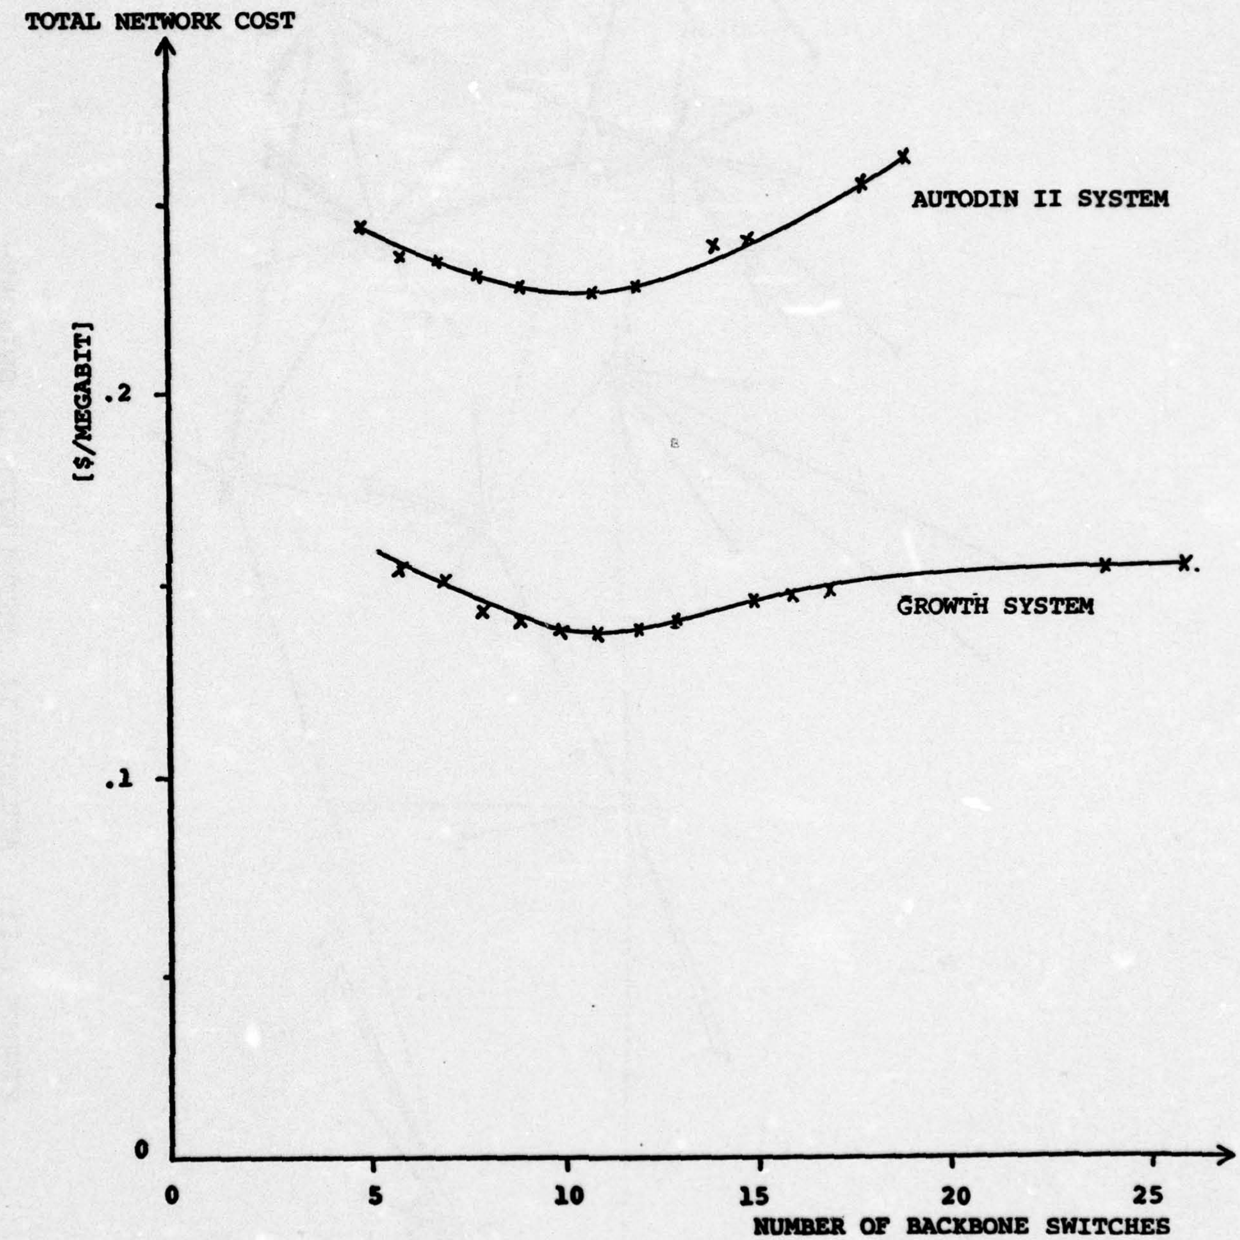

FIGURE 1.10: COST PER MEGABIT FOR AUTODIN II AND THE GROWTH SYSTEM AS A FUNCTION OF NUMBER OF BACKBONE SWITCHES

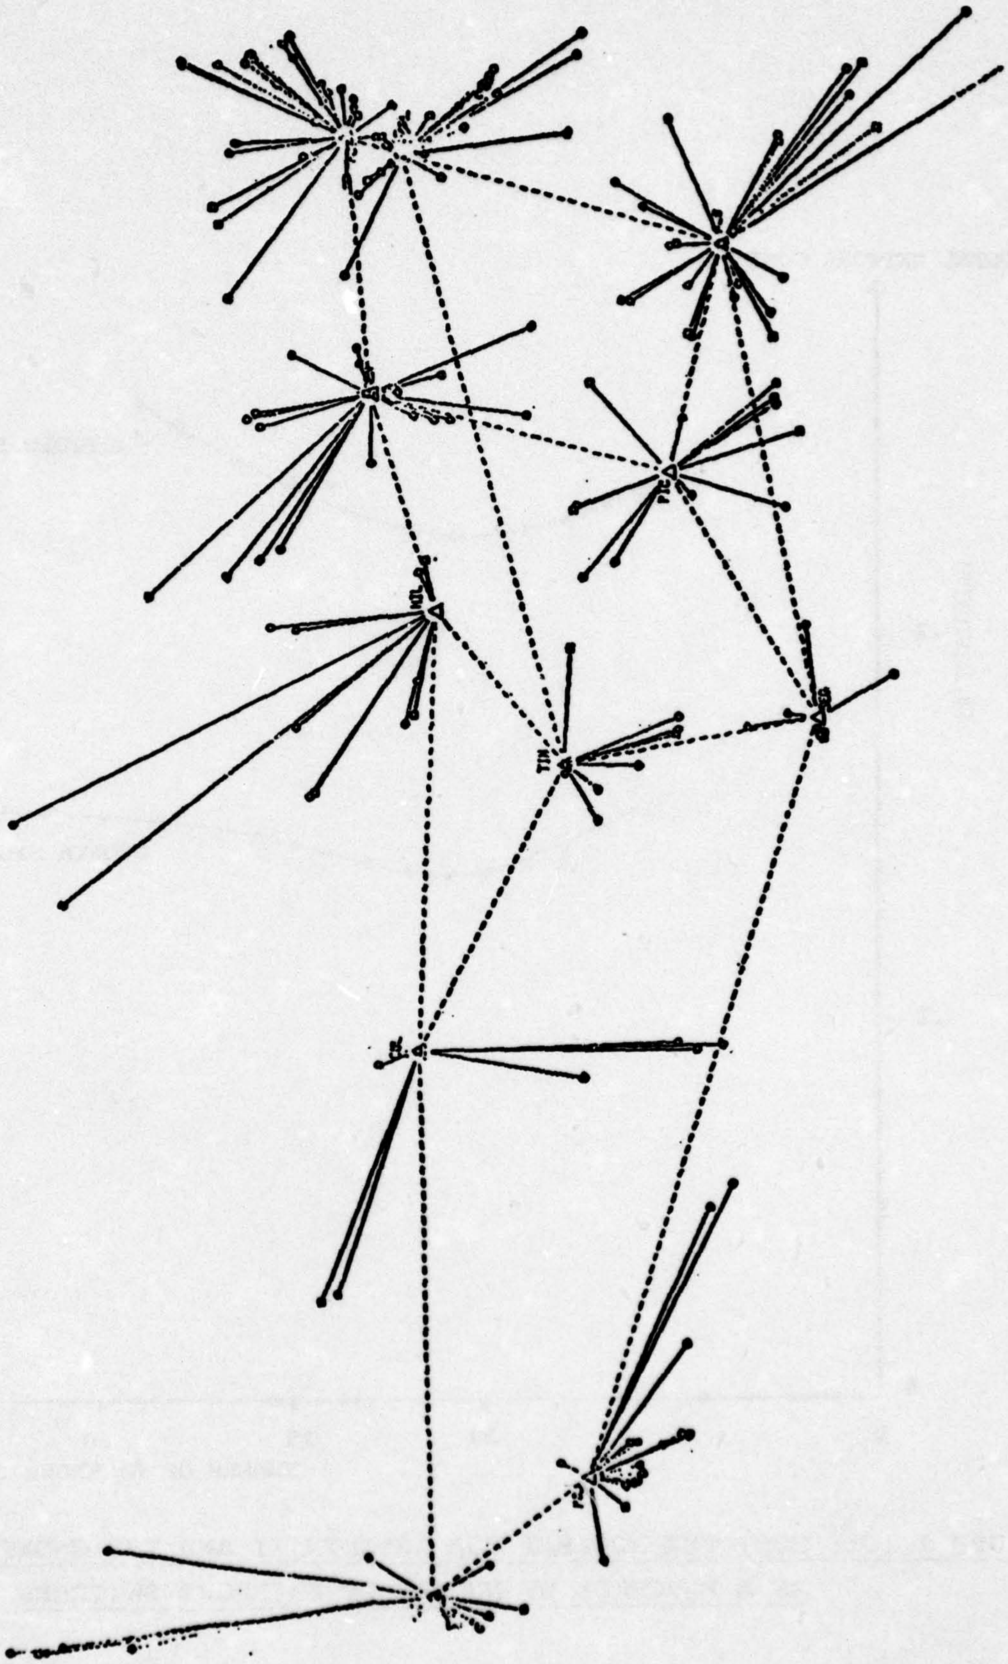

FIGURE 1.11: AUTODIN II SYSTEM WITH 11 SWITCHES

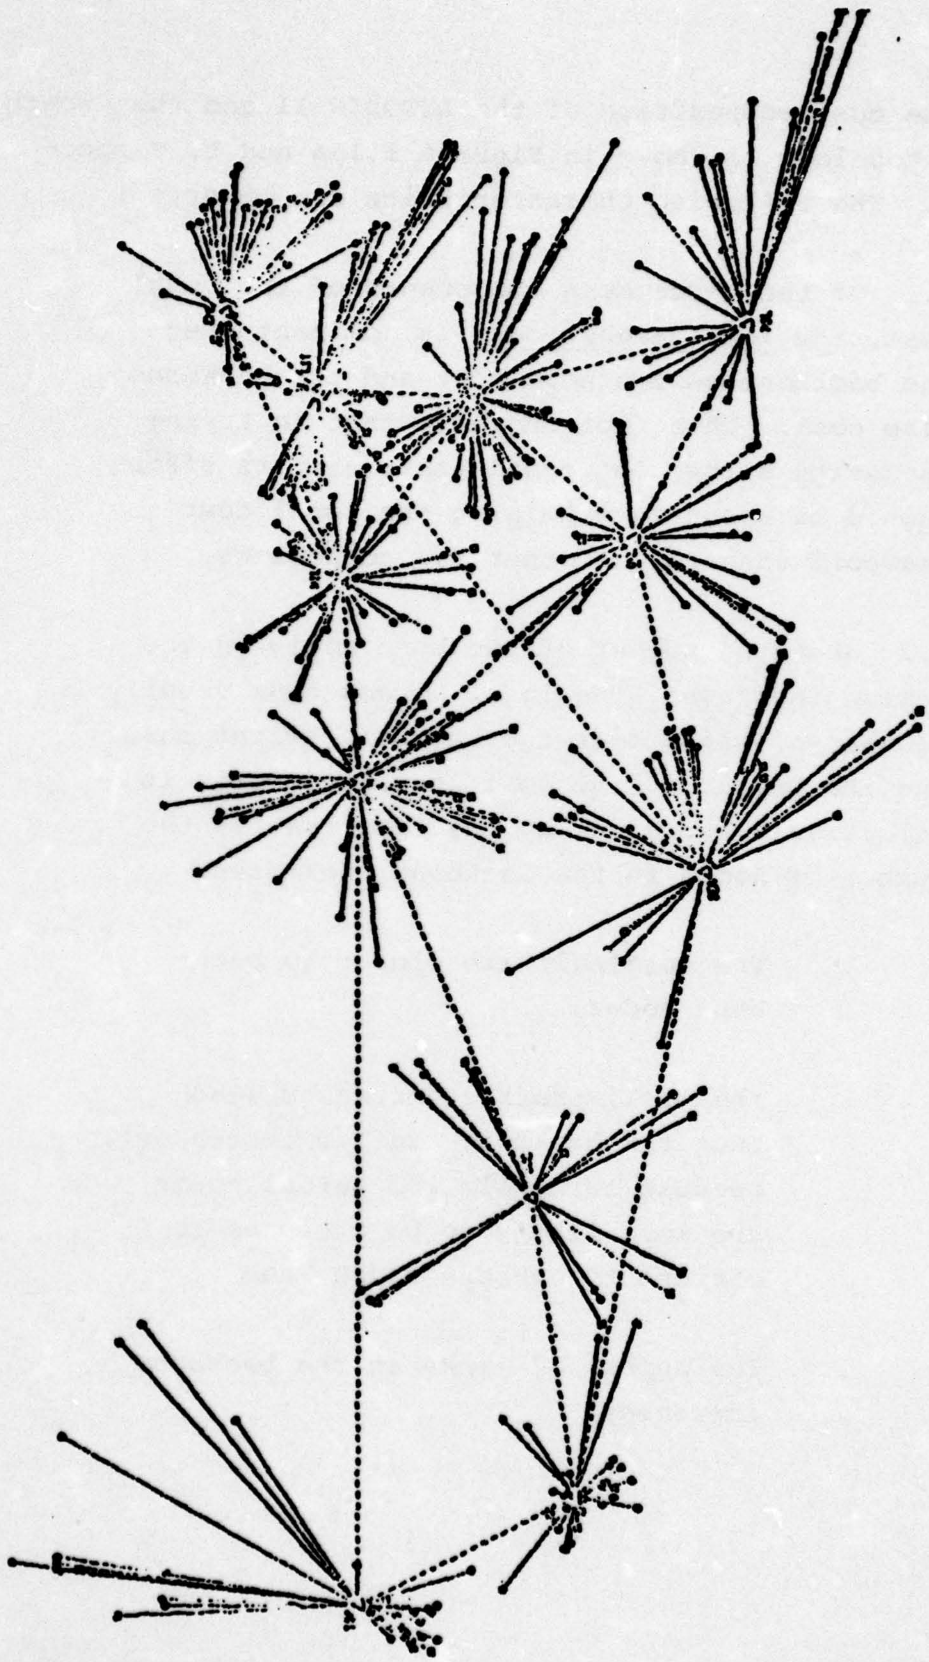

FIGURE 1.12: GROWTH SYSTEM WITH 11 SWITCHES

C. The cost composition of the AUTODIN II and the Growth System topology is shown in Figures 1.13a and b, respectively. The following characteristics are observed:

C.1 Of the three main components of the total cost, the local access cost is dominant over the backbone switch node cost and the backbone line cost. This fact suggests that in trying to design a low cost configuration, more effort should be spent in examining the local cost tradeoff than in the other two components.

C.2 When the number of backbone nodes in the design increases, the local access cost usually decreases, while both the backbone switch cost and the backbone line cost increase. This is intuitively explained by the fact that as the number of nodes in the backbone increases:

- The terminals are closer to backbone nodes,

- The local traffic decreases (and thus the backbone traffic increases), because terminals and parent Hosts are more likely to be assigned to different backbone nodes, and

- The number of links in the backbone increases.

AUTODIN II

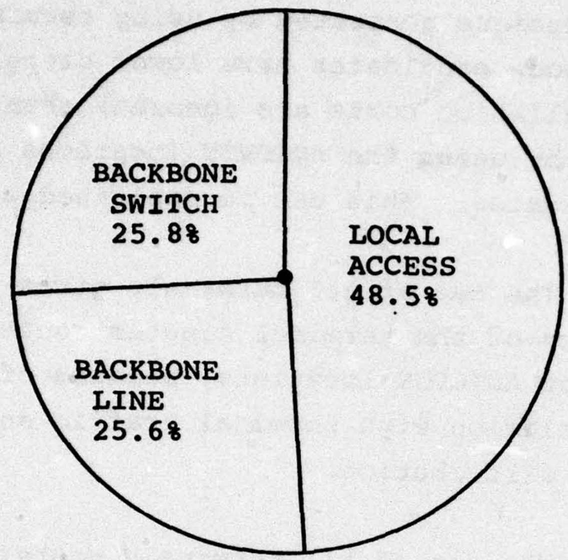

FIGURE 1.13a: COST COMPOSITION OF THE OPTIMAL BACKBONE DESIGN FOR THE AUTODIN II SYSTEMS

RGC

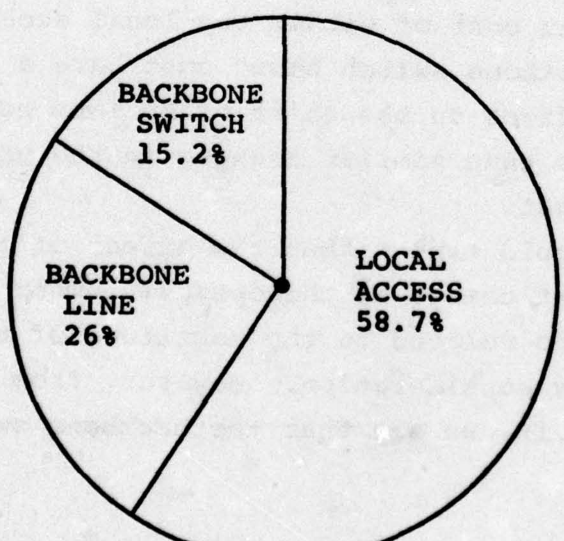

FIGURE 1.13b: COST COMPOSITION OF THE OPTIMAL BACKBONE DESIGN FOR THE GROWTH SYSTEM

D. The designs generated by using terminal sites as backbone node candidates have lower overall cost (when site installation costs are ignored) than the designs generated by using the AUTOVON locations as backbone node candidates. This can be explained as follows:

D.1 The set of all terminals gives a better choice of the terminal cluster centers than the set of AUTOVON locations, because of a better correlation with terminal traffic and geographical distribution.

D.2 The set of all terminals contains more elements than the set of AUTOVON locations, and hence offers more freedom in selecting the geographical location of switches.

E. In studying the effect of unit component cost variations on the optimum number of backbone nodes we have observed (see Figures 1.14, 1.15, and 1.16) that changes in the unit cost of either the local access mileage cost or the backbone switch basic cost have a much more pronounced effect on the shift of optimum number of backbone nodes than similar changes in the unit backbone mileage cost.

We would expect that the effect of the unit component cost change on the optimum number of backbone nodes to be related to the magnitude of the component cost under consideration. However, from Figures 1.14 through 1.16, we see that the backbone switch basic

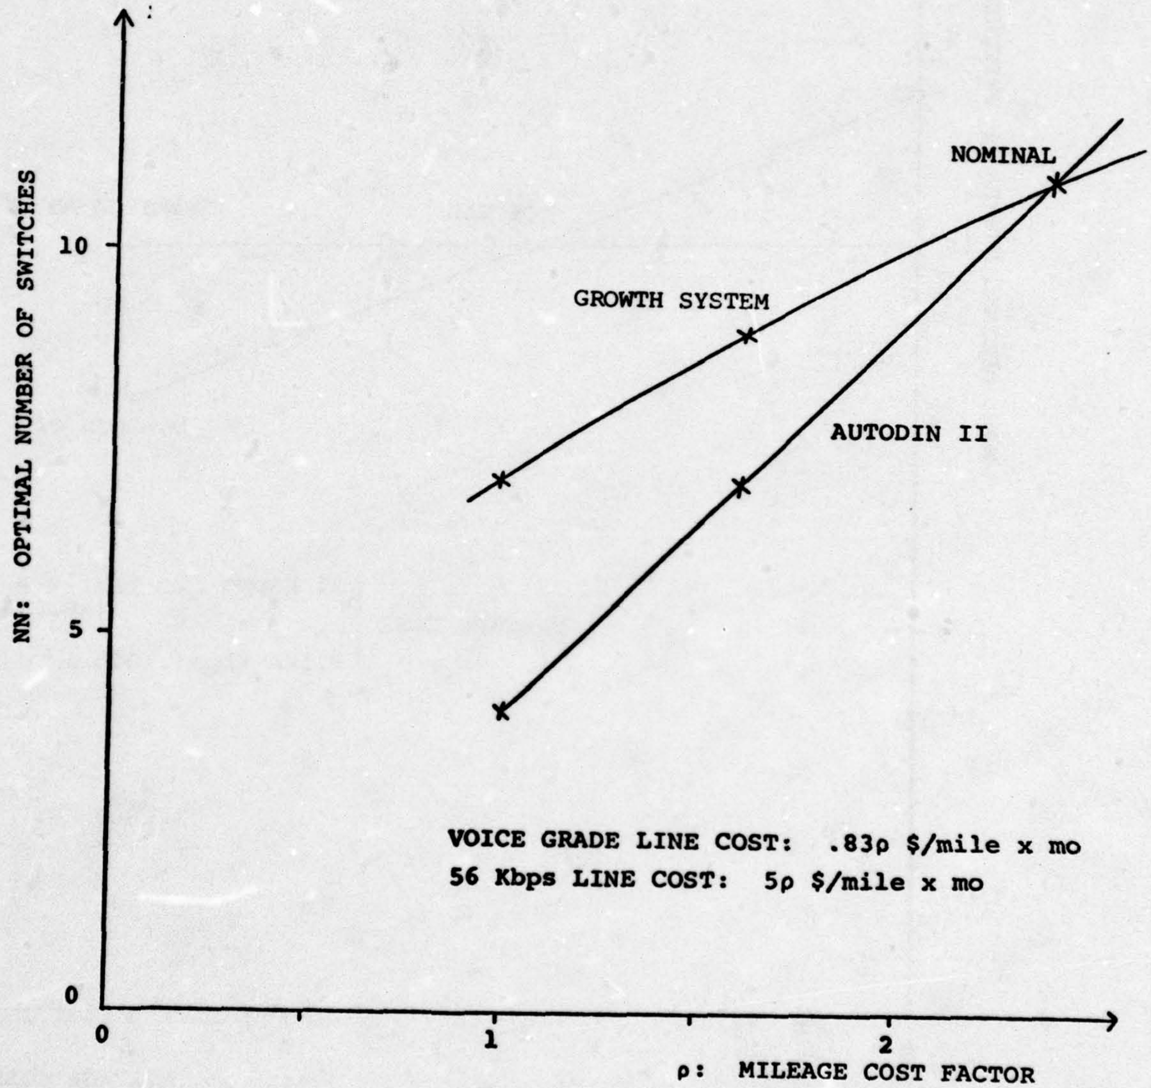

FIGURE 1.14: EFFECT OF LOCAL ACCESS MILEAGE COST VARIATIONS ON OPTIMAL NUMBER OF SWITCHES FOR AUTODIN II AND GROWTH SYSTEMS

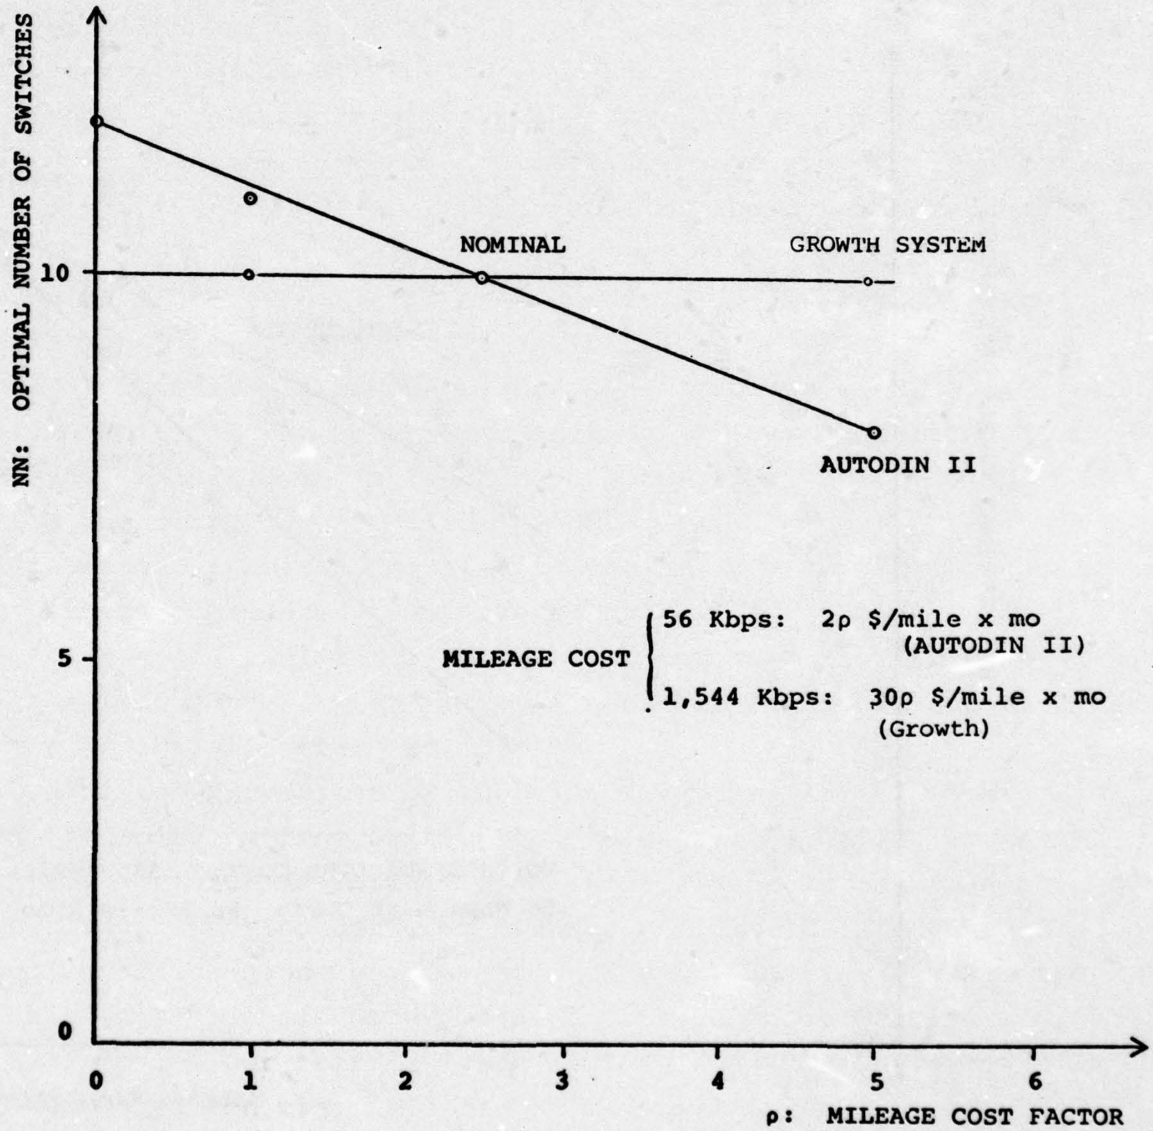

FIGURE 1.15: EFFECT OF BACKBONE TRUNK MILEAGE COST VARIATIONS ON OPTIMAL NUMBER OF SWITCHES FOR AUTODIN II AND GROWTH SYSTEMS

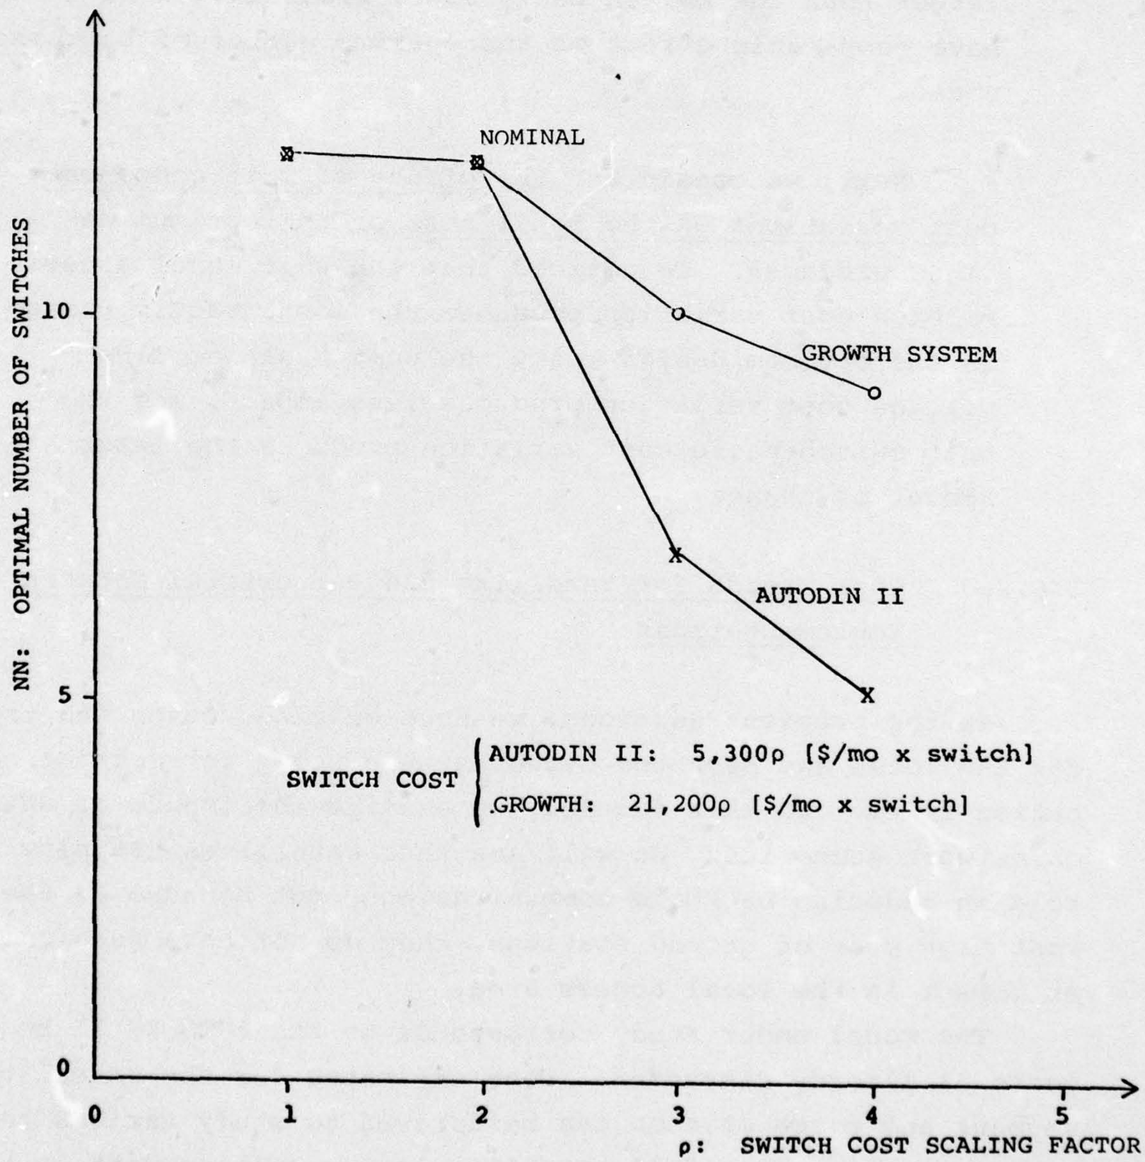

FIGURE 1.16: EFFECT OF SWITCH COST VARIATIONS ON OPTIMAL NUMBER OF SWITCHES FOR AUTODIN II AND GROWTH SYSTEMS

cost is smaller than the backbone line mileage cost, yet the unit switch basic cost variation has a stronger effect on the optimum number of backbone nodes. In fact, even though the local access line mileage cost is much larger than the switch basic cost, their unit cost changes have comparable effect on the optimum number of backbone nodes.

F. Next, we considered the effect of unit component cost variations on the total cost of the optimum designs produced. We noticed that the unit local access mileage cost variation produces the most drastic change in the optimum design cost, the unit backbone line mileage cost variation produces less impact, and the unit switch basic cost variation produces the least amount of change.

1.6.2.3 Cost Trends for Satellite and Terrestrial Network Implementations

In the previous sections, we have examined costs and tradeoffs for the local and backbone networks when using terrestrial communication lines. In this section, we examine the impact of satellites on network economics. We will see that satellites can play a major role in reducing backbone communications, but because of the current high cost of ground stations, they do not have as significant an impact in the local access area.

The model under study corresponds to the AUTODIN II requirements as already discussed. Cost estimates for the satellite space segment and earth station can be derived to study various network schemes ranging from full satellite implementation with earth stations at each site to hybrid terrestrial-satellite configurations.

The following major trends emerge:

A. When earth stations are installed at each site, the cost increases vary rapidly with the number of user sites that must be served. The cost of a terrestrial implementation, on the other hand, varies less critically with number of sites. Therefore, for a given value of throughput, there is a "threshold" value of the number of user sites above which the terrestrial network is more cost-effective, and below which the satellite network is more cost-effective. This behavior is illustrated by Figure 1.17 comparing satellite and terrestrial costs for a throughput level of 1.26 Mbps AUTODIN II requirement. In particular, for the AUTODIN II system with 300 distinct locations, the satellite implementation cost is $1,386K/mo., while the terrestrial cost is $898K/mo. The satellite network with one earth station at each site is therefore not cost-effective for AUTODIN II.

B. While an implementation with earth stations at each site may not be cost-effective, substantial savings may be obtained by installing a few stations located at appropriate centers of mass (concentration points) of the large user population. In a multilevel network structure, for example, considerable savings can be obtained by installing earth stations at each backbone switch, and replacing the backbone terrestrial network with a satellite network. End-to-end delay performance, however, may become unacceptable. As a solution, hybrid configurations may be considered consisting of a "thin"

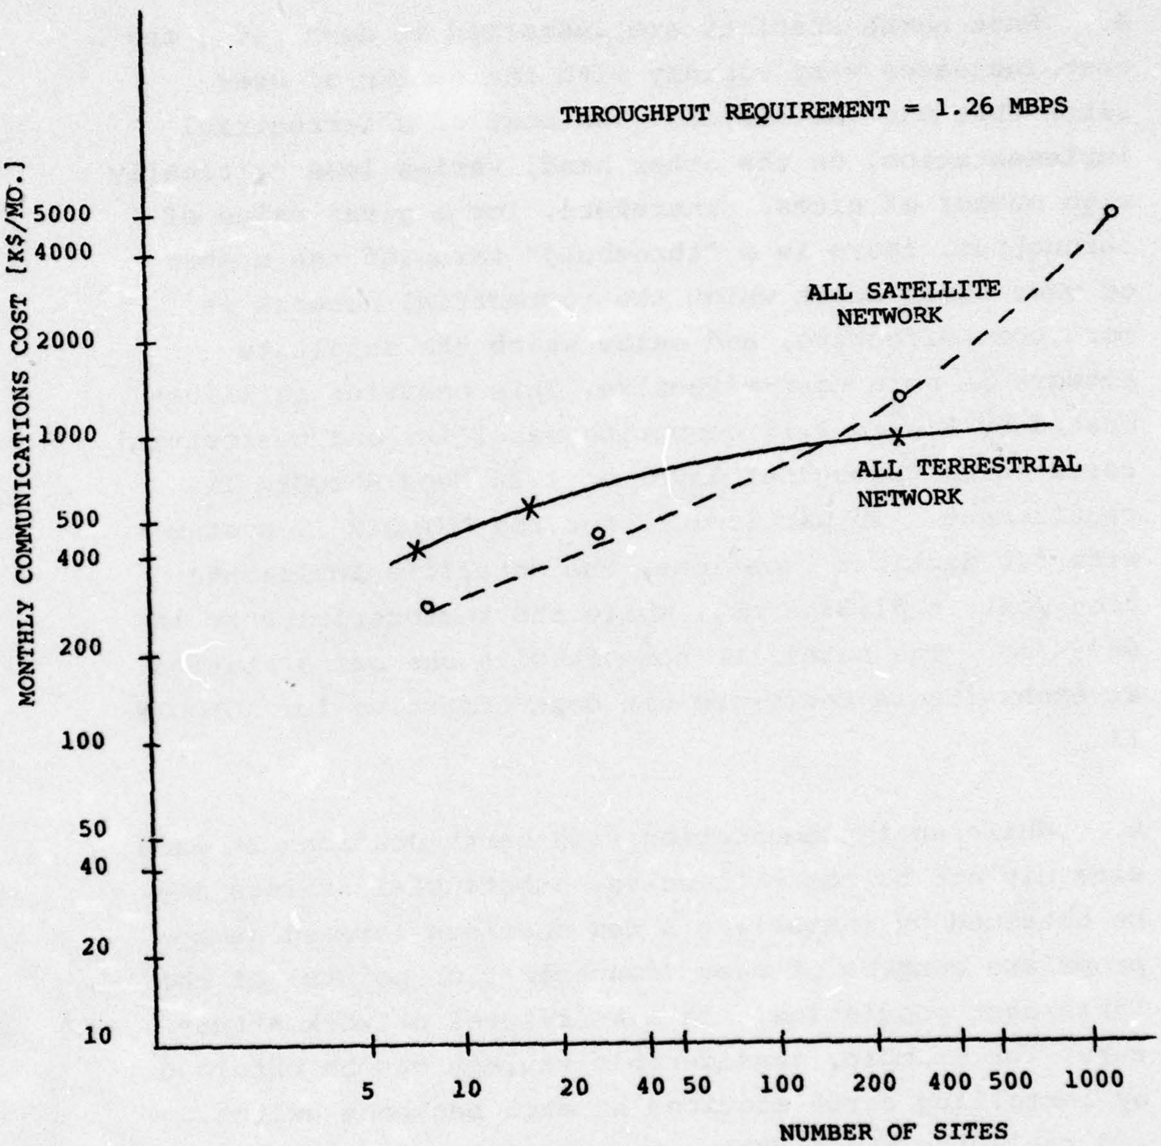

FIGURE 1.17: COMMUNICATIONS COST VERSUS NUMBER OF SITES FOR TERRESTRIAL AND SATELLITE NETWORKS

terrestrial backbone net (for low delay traffic) and a satellite backbone net superimposed on the terrestrial component (for bulk traffic). Evaluation of hybrid structures for the AUTODIN II backbone net shows that the cost of the hybrid implementation depends on the number and location of ground stations installed. In particular, the minimum cost configuration is obtained by installing ground stations at 6 out of the 8 switches of the proposed AUTODIN II. The results are summarized in Figure 1.18.

While the analysis discussed above is based on the AUTODIN II - Phase I requirements model, similar conclusions are reached when the Growth System with an order of magnitude more traffic is examined. In this case, the savings generated by using satellite channels within the backbone network are even more substantial while the impact in the local area remains the same.

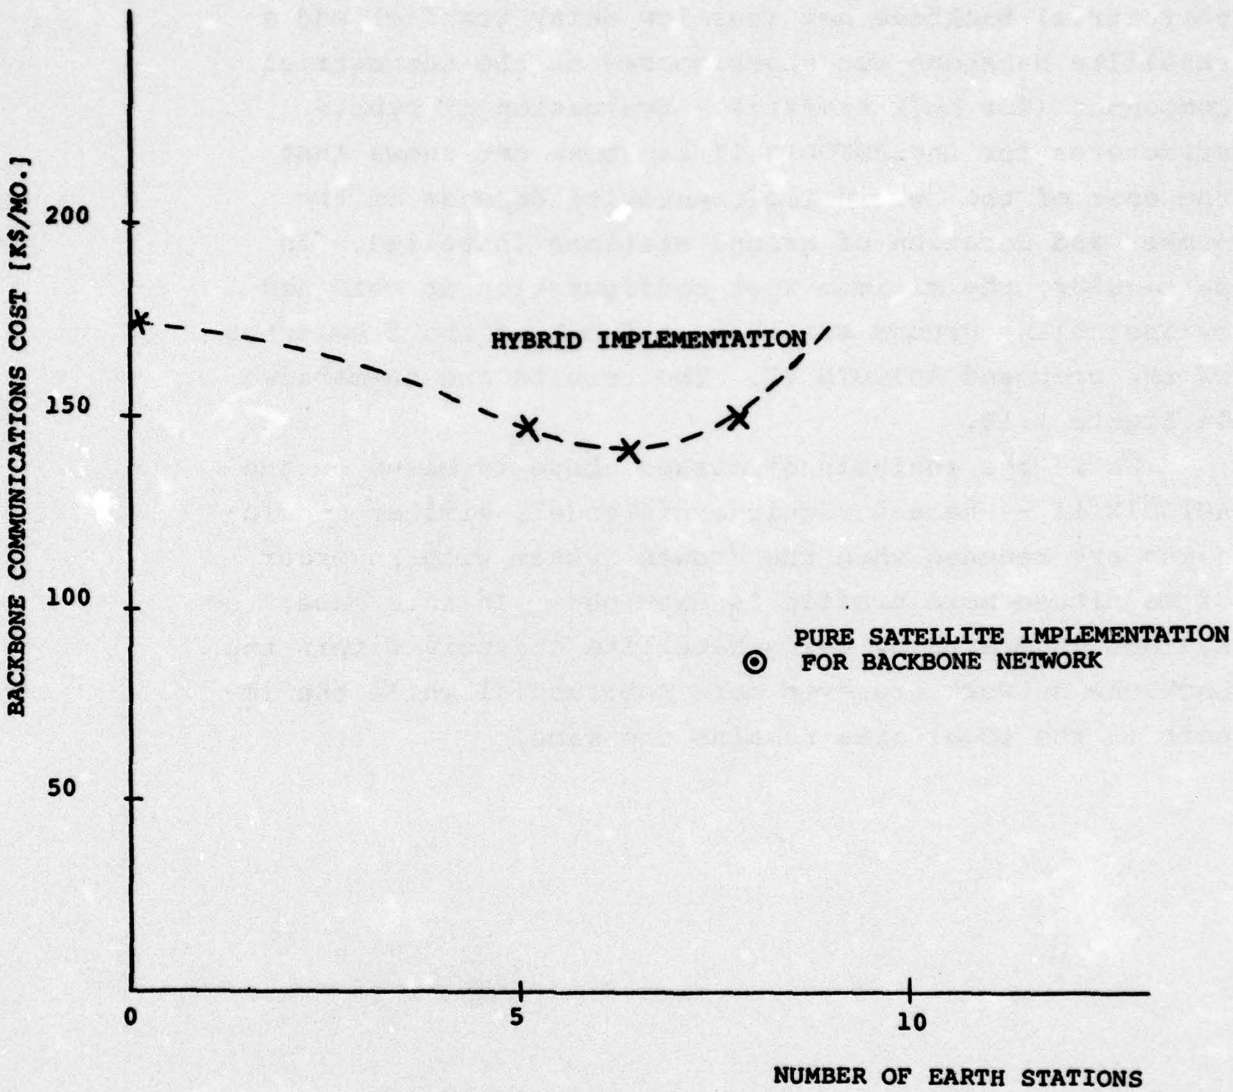

FIGURE 1.18: BACKBONE COMMUNICATIONS COST (EXCLUDING SWITCH COST) FOR A HYBRID TERRESTRIAL - SATELLITE IMPLEMENTATION, AS A FUNCTION OF NUMBER OF EARTH STATIONS

1.7 CONCLUSIONS

The studies we have performed are not exhaustive since we have not considered issues such as security, management, operations and control of the network. However, within the area of the subjects discussed, we feel the results are representative and indicative of the cost and performance trends that will appear when these large scale networks are built. In addition, the results identify several emerging technologies (e.g., packet radio for conventional local access for data) which can meet DOD network requirements at lower cost than conventional techniques.

This report has summarized the results of investigations into local, regional, and large scale packet switched network cost and performance for Defense Department communications. The issues examined include:

- Alternative large network architectures.
- Performance evaluation criteria.
- Formulation of the large network design problem.
- Defense communication requirements.
- Local access network tradeoffs.
- Backbone and total network performance and cost characteristics.
- Applications to AUTODIN II.

Conclusions of the effort include:

- In a large, well-designed network, the most cost-effective network architecture is a multilevel hierarchical structure. It contains a high level backbone network and low level local access networks. Both the high and low level networks may be hierarchical. For example, the backbone network may contain satellite and terrestrial subnetworks.

- In a large optimized hierarchical network, delay and reliability in the backbone network are typically one order of magnitude better than in the local distribution network (except for satellite delays) and therefore variations in backbone delay and reliability requirements have little effect on global performance.

- Local access costs are on the order of 50% (or more) of total network costs. Moreover, total system cost is most sensitive to changes in local access component cost variations. Therefore, the local distribution network should receive high priority for cost minimization.

- Switch hardware and backbone communication line costs share the remaining cost on about an equal basis. This continues the cost trends established by ARPANET.

- The most cost-effective <u>conventional</u> local distribution alternative consists of dedicated point-to-point lines which achieve cost savings via clustering of colocated terminals and Hosts by means of multiplexing or concentration.

- Multidrop polling techniques on tree or loop structures do not offer sufficient savings to warrant the changes required to overcome implementation difficulties and security shortcomings.

- Satellite use for local access via ground stations located at terminal or Host facilities does not appear to be cost-effective at current ground station costs and kilobit data rate requirements per facility.

- Packet radio provides a cost-effective alternative to conventional local access techniques. Packet Radio systems can meet the high reliability requirements for critical DOD systems. If PRU's were manufactured in large quantities for $20,000 - $30,000 or less, PRNET designs would be more cost-effective than conventional designs. These costs are in the range of current technology and trends indicate that these units can be built for substantially less than this cost.

- Local access and backbone network cost and performance is interrelated by the number, location, and capacity of the backbone switching nodes. Total system cost versus the number of backbone nodes does not vary significantly over a wide range of backbone node numbers and hence, the decision to distribute switches to many locations or consolidate switching as in AUTODIN II is relatively independent of communication network economics.

- The optimum number of backbone nodes for AUTODIN II - Phase I ranges from 9 to 12 sites provided that switching facilities have sufficient capacity to handle traffic requirements. However, cost differentials are less than 10% when the number of backbone nodes are increased substantially. Hence, wide distribution of switches, a desirable survivability feature, may be achieved at small incremental hardware and line cost.

- Considerable backbone line cost savings (on the order of 50%) can be achieved by installing satellite ground stations at all backbone nodes and eliminating all terrestrial backbone links. However, in this case, end-to-end delay may become unacceptable.

- A hybrid satellite/terrestrial backbone network is cost-effective and provides acceptable delay performance. In this case, the minimum cost configuration for AUTODIN II involves installing ground stations at 6 out of the 8 proposed AUTODIN II switch locations. This results in a savings of about 25% of the backbone line costs.

- Networks with an order of magnitude more traffic than AUTODIN II can be implemented using the same basic architecture and concepts as the AUTODIN II system and can provide the required delay, reliability, and throughput. Such networks can achieve very high reliability and can be operated at lower costs per unit of traffic than smaller networks.

- Communication costs per megabit of transmitted data in AUTODIN II should range between 23-25 cents. This cost level is comparable to costs projected by NAC for ARPANET and similar networks as early as 1972.

- Economies of scale, which have previously been demonstrated within ARPANET and other medium size packet networks exist within the large network DOD environment. A tenfold increase in traffic from the AUTODIN II 1.26 Mbs requirement results in about a 45% decrease in cost per transmitted megabit of data (to a cost on the order of 14 cents per megabit).

- It is the opinion of the authors that this study and past efforts demonstrate that packet switching can provide cost-effective, reliable, and responsive performance for DOD data networking. Furthermore, the benefits of packet switching, first demonstrated for ARPANET, extend to AUTODIN II - Phase I and to networks with an order of magnitude more traffic than AUTODIN II - Phase I. Thus, approaches adopted today for current DOD requirements can be expected to survive the addition of new requirements.

VOLUME 1

REFERENCES

REFERENCES

[FRANK, 1976] Frank, H., I. Frisch, M. Gerla, and I. Gitman, "A Comparison of Local Access Techniques for Packet Switched Networks", Proceedings of Electro 1976, May, 11-15, 1976, Boston, Mass.

[FRISCH, 1975] Frisch, I. T., "Technical Problems in Nationwide Networking and Interconnection," IEEE Transactions on Communications, Vol. COM. 23, No. 1, January, 1975, pp. 78-88.

[KAHN, 1975] Kahn, R. E., "The Organization of Computer Resources Into A Packet Radio Network", Proceedings of the National Computer Conference May, 1975, pp. 233-243.

[NAC, 1973] Network Analysis Corporation, The Practical Impact of Recent Computer Advances on the Analysis and Design of Large Scale Networks, First Semiannual Technical Report, May, 1973, available from Defense Communication Center, Arlington, Va.

[NAC, 1973A] Network Analysis Corporation, The Practical Impact of Recent Computer Advances on the Analysis and Design of Large Scale Networks, Second Semiannual Technical Report, December, 1973, available from Defense Communication Center, Arlington, Va.

REFERENCES (Cont'd.)

[NAC, 1974] Network Analysis Corporation, The Practical Impact of Recent Computer Advances on the Analysis and Design of Large Scale Networks, Third Semiannual Technical Report, June, 1974, available from Defense Communication Center, Arlington, Va.

[NAC, 1974A] Network Analysis Corporation, Computer Network Requirement Study, March, 1974.

[NAC, 1975] Network Analysis Corporation, Alternative Network Strategies for Defense ADP Communications, July, 1975.

[NAC, 1976] Network Analysis Corporation, Local and Regional Communication Networks - Technologies and Architectures, Sixth Semiannual Technical Report, Vol. 4, February, 1976.

VOLUME 1

APPENDICES

APPENDIX

EXECUTIVE SUMMARY

TO

ALTERNATIVE NETWORK STRATEGIES

FOR

DEFENSE ADP COMMUNICATIONS

JULY, 1975

The basic cost data that was used to compute the results of this study was based on preliminary estimates. The results therefore should be interpreted as indicative rather than reflecting exact costs.

Sponsored by

ADVANCED RESEARCH PROJECTS AGENCY
Department of Defense

Contract No. DAHC 15-73-C-0135

NAC REPORT NO. TR.048.01

Contractor

NETWORK ANALYSIS CORPORATION
Beechwood, Old Tappan Road
Glen Cove, New York 11542

ALTERNATIVE NETWORK STRATEGIES FOR DEFENSE ADP COMMUNICATIONS

EXECUTIVE SUMMARY

OBJECTIVE

This report presents the results of a cost-evaluation study of alternative Defense Communication network strategies. In this study, the terminal and traffic requirements of 35 ADP systems are considered, and the problems of designing the data networks which accommodate these requirements at minimum cost is addressed. The goal of the study is to identify the communications line and hardware costs associated with the array of feasible methods for implementing each system requirement.

SYSTEM OPTIONS CONSIDERED

The system options considered can be placed in two categories: systems without switching and integrated switched systems. Alternatives examined in each category are described below.

Systems Without Switching

1. Separate systems with terminals and hosts on dedicated lines. Such systems require no overall network optimization. For example, whenever new terminals are added to a system, separate dedicated lines are leased to the appropriate hosts.

2. Separate systems with local line sharing. This approach requires a manager at each facility to coordinate the use of facility multiplexers and/or concentrators and the ordering of lines from a given location to reduce communication cost.

3. Separate systems with local and regional line sharing. This approach requires a manager for each system who is responsible for optimizing each system network design.

4. Limited system integration without switching. In this approach, several systems may share (via multiplexers) high speed lines. A coordinator with the ability to configure each of the systems and to combine requirements where appropriate is required.

<u>Fully Integrated Packet-Switched Systems</u>

The integrated packet-switched network approach allows the joint use of communication lines, communication equipment and host computers for resource and load sharing applications. Designs for three basic network approaches are developed.

1. A fully integrated network with packet switches located at eight AUTODIN I store-and-forward switch sites.

2. A fully integrated network in which switch locations are selected to minimize the overall communication line plus hardware cost without regard to the specific security issues at each site.

3. Two independent packet-switched networks, one handling only the encrypted traffic generated by the twelve systems requiring encryption and the other handling all remaining traffic.

Each of the above system approaches requires an overall systems manager with knowledge of the individual system requirements and full control of the access and backbone network designs.

For each of the above alternatives, the cost impact of two different packet switch alternatives is assessed. These are:

1. Use of "Pluribus" high speed, modular, multiprocessor IMP's currently under development.

2. Clustering of currently available ARPANET IMP's at switching facilities to meet high bandwidth traffic requirements.

Also examined for each network strategy are link by link and end-to-end encryption alternatives.

ASSUMPTIONS

Assumptions include traffic volumes to be accommodated, component costs, message response times, reliability, and security requirements. The approach taken is to design each alternative so that it meets all system requirements. The costs associated with each alternative are then compared. A summary of the assumptions used in the study is given below.

System Size

1.	Number of host computers:	87
2.	Number of terminals:	1,103
3.	Total traffic:	1.26 Megabits/second

Data rates used for the study are in some cases best guess estimates that cannot, at present, be verified.

Cost Factors

Cost factors are based on current procurement estimates for tariffed communication lines and hardware. Communication line costs include mileage, termination and modem charges. Hardware cost factors include purchase price, installation, initial support, operations, maintenance, and amortization. Cost factors not considered are the host processor cycle time costs required to support various network connection schemes, network management costs, and the security costs of specially cleared switches and operating personnel required for link encrypted alternatives.

Reliability

1. Availability greater than or equal to .99 for non-critical systems.

2. Availability greater than or equal to .9995 for critical systems.

3. The critical systems contain 34 host computers, 120 terminals, and 12.4% of the total system traffic.

Message Response Time

1. For the independent, non-switched systems, the average end-to-end delay for a 500 bit transmission between terminal and host or host and host must be less than 1 second.

2. For the integrated systems, the average delay for a 500 bit transmission must be:

- Less than 1 second between terminal and backbone switch

- Less than .25 second between host and backbone switch

- Less than .1 second between any pair of backbone switches

The impact on response time and bandwidth of end-to-end, user specific protocols was not examined.

Security

1. 12 of the 35 systems require encryption.

2. The 12 systems requiring encryption contain 56% of the total number of hosts, 6% of the total terminals, and generate 33% of the total network traffic.

CONCLUSIONS

Systems Without Switching

1. More than seven million dollars per year can be saved, when compared to independent systems with hosts and terminals on dedicated lines, by introducing multiplexing at the facility level. This implies that each facility must have a manager to promote line and hardware sharing for the terminals and hosts at that location.

2. An additional one million dollars per year can be saved by introducing regional multiplexing (or concentration) within each individual system. This implies that an overall network manager for each system is required to optimize the design of each configuration.

3. If each system is separately optimized, little additional advantages are achieved by combining systems without adding intersystem switching.

Integrated Packet Switched Systems

1. A preliminary packet-switched integrated network design yields a total system whose cost is within 13% of the best non-switched alternatives. This approach requires an overall system manager with control over the backbone and local access network configuration.

2. Additional savings, using alternatives such as domestic satellites and new IMP minicomputers not studied in detail in this report, are achievable in a fully integrated network.

Thus, resource and load sharing capabilities, inherent in a fully integrated system, can be achieved at no more than a small incremental cost when compared to the best non-switched system.

3. Savings of over seven million dollars per year are achieved via an integrated network when compared to the strategy of independent systems with dedicated host and terminal communication lines.

4. An alternative to handling all messages from all systems on a single integrated network is to construct two separate networks, one handling the messages which require encryption, the other handling the remaining traffic. A two system approach was developed with the secure system having eight backbone switching nodes and with the system handling unencrypted traffic having 27 backbone nodes. The total cost of the two systems is approximately 7% higher than the cost of a single integrated system.

5. If a two system approach were adopted, communications for current ARPANET users could be provided through the 27 backbone node system accommodating the unencrypted traffic. A feasible strategy would be to replace ARPANET IMP's by Very Distant Host Interfaces (VDH's) connected to ports on IMP's at military bases. VDH hardware plus access lines would cost less than one million dollars per year. The incremental backbone network cost to handle ARPANET traffic would be nominal. (As a point of comparison, ARPANET line cost alone will exceed 1.5 million dollars per year in the near future.)

6. End-to-end encryption is more cost-effective than link encryption for all non-switched alternatives, even when the cost of the secure switches and operating personnel required for link encryption is ignored. For the switched systems, current annual costs for end-to-end encryption would range from $12,000 to $300,000 (1% to 45%) more than link encryption, when the switch factor is ignored. Since even gross estimates of the switch factor cost greatly exceed this difference, end-to-end encryption is the most cost-effective alternative.

7. The number of switch locations in a single integrated network does not strongly influence the sum of the communication and hardware costs over the range of locations considered. However, the number of locations has a major influence on the strategy used to implement the packet switches. In addition, systems with few switch sites require a large number of Very Distant Host interfaces which may create excessive host CPU overhead.

8. The major element of technological risk in the switched system approach is the development of the backbone switching facilities. If the integrated system contains a small number of backbone switches (e.g., eight), each switch must have a throughput capacity many times greater than that of the current ARPANET IMP. The highest risk strategy involves the development of high speed Pluribus IMP's.

9. A low risk alternative is the modification of ARPANET IMP's by software changes and core expansion to handle DOD priority and preemption requirements. These IMP's would then be interconnected in "clusters" at the switch facility locations to handle the required traffic load. If a small

number of backbone sites are utilized, then the IMP clusters at each site may contain as many as 16 IMP's, but are capable of handling the projected traffic requirements. While this approach appears to be "inelegant," it is shown to be feasible and the cost basis for comparing alternatives includes factors such as floor space and other major considerations.

10. If IMP's are distributed to a larger number of facilities, the IMP clusters are considerably smaller. Thus, a 27-switch site system requires on the average four IMP's per site.

11. The costs associated with the high speed Pluribus IMP and the modification of current ARPANET IMP's are approximately equal. However, the risks are substantially different.

12. A third alternative, of intermediate risk, is to develop a second generation ARPANET IMP based on currently available proven minicomputers. This approach would require ARPANET software modifications for compatibility as well as for the addition of the required new software capabilities. The use of a new generation of ARPANET IMP's would reduce the IMP cost by at least a factor of two.

13. If a new ARPANET IMP were developed, the 27-switch site system would be the most economical of all of the integrated packet-switched approaches examined, with a cost approximately 9% less than the comparable eight site system.

14. If a new ARPANET IMP were developed, the two system approach, which segregates encrypted and unencrypted messages on different networks would be approximately equal in cost to that of the single integrated network with Pluribus message processors.

15. The sensitivity of the study results to large increases in traffic over those used has not been examined. However, in this case, a major limiting factor would then be switch capacity. If, for example, switches were restricted to eight locations, the only way of handling the traffic would be to create clusters of high speed IMP's. A more flexible long range strategy would be to distribute packet switches to larger numbers of locations to meet long range growth requirements.

16. Additional considerations for full comparison of alternatives include the management and organizational issues involved in implementing a fully distributed integrated network.

UNCLASSIFIED
Security Classification

DOCUMENT CONTROL DATA - R&D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) NETWORK ANALYSIS CORPORATION BEECHWOOD, OLD TAPPAN ROAD GLEN COVE, NEW YORK 11542	2a. REPORT SECURITY CLASSIFICATION Unclassified 2b. GROUP None

3. REPORT TITLE
SIXTH SEMIANNUAL TECHNICAL REPORT, FEBRUARY 1976, FOR THE PROJECT, "LOCAL, REGIONAL AND LARGE SCALE INTEGRATED NETWORKS," (VOLUMES 1 THROUGH 5)

4. DESCRIPTIVE NOTES (Type of report and inclusive dates)
SIXTH SEMINANNUAL REPORT, FEBRUARY 1976

5. AUTHOR(S) (Last name, first name, initial)
NETWORK ANALYSIS CORPORATION
BEECHWOOD, OLD TAPPAN ROAD
GLEN COVE, NEW YORK 11542

6. REPORT DATE	7a. TOTAL NO. OF PAGES	7b. NO. OF REFS
FEBRUARY 1976	880	140

8a. CONTRACT OR GRANT NO. DAHC-15-73-C0135 b. PROJECT NO. c. ARPA ORDER NO. 2286 d.	9a. ORIGINATOR'S REPORT NUMBER(S) SEMIANNUAL REPORT 6 (5 VOLUMES) 9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report) ---

10. AVAILABILITY/LIMITATION NOTICES

This document has been approved for public release and sale; its distribution is unlimited.

11. SUPPLEMENTARY NOTES	12. SPONSORING MILITARY ACTIVITY
None	Advanced Research Projects Agency, Department of Defense

13. ABSTRACT This technical report presents significant accomplishments in the development of computational techniques for network analysis and design, and the following cost/performance tradeoffs for data and computer communication networks: cost-effectiveness and economies of scale previously demonstrated for medium size packet switching networks (e.g., ARPANET) extend to large scale networks (e.g., 10 Mbs throughput). The most cost-effective architecture for large scale networks is a multi-level hierarchical structure. The cost of hierarchical networks is not very sensitive to number of backbone nodes near the optimum; the local access cost is approximately 50% of the total, and is the most sensitive component. For AUTODIN II, Phase I, the minimum cost backbone network configuration is a hybrid terrestrial-satellite net with six (out of eight nodes) ground stations, resulting in 25% savings in backbone line cost compared to a terrestrial net only. Packet Radio local access is a cost-effective alternative to conventional techniques, providing a Packet Radio Unit can be manufactured for $20K-$30K. Significant results are reported on Packet Radio network communication protocols, throughput-delay-reliability analysis, and stability control.

14. KEY WORDS Computer networks, communication networks, terrestrial and satellite networks, packet radio networks, throughput, cost, delay, blocking, ARPA Computer Network, store-and-forward, packet switching.

UNCLASSIFIED
Security Classification